Intel Express 8100 Router

Reference Manual

Second edition April 1998

Copyright © 1998, Intel Corporation. All rights reserved.

Intel Corporation, 5200 NE Elam Young Parkway, Hillsboro, OR 97124-6497

Intel Corporation assumes no responsibility for errors or omissions in this manual. Nor does Intel make any commitment to update the information contained herein.

* Other product and corporate names may be trademarks of other companies and are used only for explanation and to the owners' benefit, without intent to infringe.

Contents

Part I LAN and WAN Links and Services

1 LAN and WAN Services in the Router 1
LAN Services 1
WAN Services 2
WAN Services and Protocols Available 3
General Facilities for WAN Services 3

2 Leased Lines Links 8
Point-to-Point Protocol (PPP) 8
Data Compression 11
Data Encryption 11
Peer Authentication using the Challenge Handshake Authentication Protocol (CHAP) 12
Peer Authentication using the Password Authentication Protocol (PAP) 15

3 Frame Relay Services 17

4 ISDN Services 20
Integrated Services Digital Network (ISDN) in the Intel Express 8100 Router 20
ISDN Concept 21
ISDN Services 22
ISDN Numbering and Addressing 24
BRA and PRA Numbering 26
MSN (Multiple Subscriber Number) 27
SUB (Sub-addressing) 27
DDI (Direct Dialing In) 27
Permanent ISDN 28
ISDN Standards 28
ISDN Network Interface Reference Configurations 30

5 X.25 Services 32
X.25 Characteristics 32
X.25 Services and Features 34
X.25 Packets and Virtual Calls 35
X.25 Addressing 37

6 PPP Multilinks 39

7 Internet Tunnels 41

Part II IP, Novell IPX and Bridging Services

8 IP Routing 47
IP Concept 47
IP Addresses Structure 48
IP Subnets 49
Address Resolution 51
Internet Control Messages 52
IP Filters—Firewall Protection 52
IP Filter Example 55
Network Address Translation (NAT) 57
Static Mapping 60
Dynamic Mapping 60
Examples of Use 62
Routing Information Protocol (RIP) 64
RIP-1 and RIP-2 64
RIP Metrics 66
Triggered RIP 67
Static Routes 68
Static Routes Example 70

9 Novell IPX Routing 71
Novell Routing Concept 71
Internetworking Packet Exchange (IPX) 72
Frame Types 72
Novell Addressing 72
Novell Metrics 73
Data Transmission 75
IPX over WAN Links 76
IPX WAN Protocol 76
IPX Watchdog Packets 77
IPX Serialization Packets 79
Sequenced Packet Exchange (SPX) 79
IPX Data Filters 82
IPX Filters Example 84
Routing Information Protocol (RIP) 85
RIP Filtering 86
Service Advertising Protocol (SAP) 87
SAP Filtering 87
Static Routes and Services 88
Configuration Example—Cost-Reduced WAN 91

10 WAN Bridging Services 93
Introduction to WAN Bridging 93
How Bridging Works 94
Forwarding Tables 95
Operation of the Bridge in the Intel Express 8100 Router 101
Controlled Bridging 106

Spanning Tree Algorithm 108
Bridging Loops 108
Spanning Tree Concept 109
Spanning Tree in the Intel Express 8100 Router 111
Bridge Spoofing 114
BPDU Spoofing 114
Non-BPDU Spoofing 115
Enabling Bridge Spoofing 115

Appendices

A Protocol Bandwidth Overheads and Requirements 117

B TCP and UDP Port Numbers 119

C Ethernet Type Codes 121

D IP Frame Formats 126

E Novell IPX Frame Formats 127

F Novell IPX Service Types 130

Preface

Products covered

The following products are covered in this manual:

- Express 8100 Router with Frame Relay
- Express 8100 Router with X.25
- Express 8100 Router with an ISDN S/T port
- Express 8100 Router with an ISDN U port

Acronyms

BACP	Bandwidth Allocation Protocol
BCP	Bridging Control Protocol
BECN	Backward Explicit Congestion Notification
BPDU	Bridging Protocol Data Units
BRA	Basic Rate Access
BRI	Basic Rate Interface
CBC	Cypher Block Chaining
CCITT	Comité Consultatif International Télégraphique et Téléphonique
CCP	Compression Control Protocol
CHAP	Challenge Handshake Authentication protocol
CRC	Cyclic Redundancy Check
DCE	Data Communication-terminating Equipment
DE	Discard Eligibility
DLCI	Data Link Connection Identifier
DLCMI	Data Link Connection Management Interface (also called LMI)
DN	Directory Number
DTE	Data Terminal Equipment
DVMRP	Distance Vector Multicast Routing Protocol
ECP	Encryption Control Protocol
FECN	Forward Explicit Congestion Notification
FR	Frame Relay
HDLC	High-level Data Link Control
ICMP	Internet Control Message Protocol
IGMP	Internet Group Management Protocol
IP	Internet Protocol
IPCP	Internet Protocol Control Protocol
IPX	Internetwork Packet Exchange
IPXCP	Internetwork Packet Exchange Control Protocol
ISDN	Integrated Services Digital Network
ITU	International Telecommunication Union
LAN	Local Area Network
LAPB	Link Access Procedure Balanced
LAPF	Link Access Procedure for Frame Relay
LCP	Link Control Protocol
LSP	Link State PDU
MIB	Management Information Base

NAT	Network Address Translation
NCP	Network Control Protocol
PAP	Password Authentication Protocol
PING	Packet InterNet Groper Function
PDN	Public Data Network
PRA	Primary Rate Access
PRI	Primary Rate Interface
PSDN	Packet Switched Data Network
PVC	Permanent Virtual Circuit
RIP	Routing Information Protocol
RSVP	ReSerVation Protocol
SAP	Service Advertising Protocol
SLIP	Serial Link Internet Protocol
SNAP	SubNetwork Access Protocol
SNMP	Simple Network Management Protocol
SPX	Sequenced Packet Exchange
SVC	Switched Virtual Circuit
TCP	Transmission Control Protocol
TFTP	Trivial File Transfer Protocol
UDP	User Data Protocol
WAN	Wide Area Network
WWW	World Wide Web

Related documents

Novell NetWare Link Services Protocol Specification Rev 1.0.
Novell part No. 100-001708-002.

IPX Router Specification.
Novell part No. 107-000029-001

User Datagram Protocol (UDP).
RFC 768

Trivial File Transfer Protocol (TFTP) revision 2.
RFC 783

Internet Protocol (IP).
RFC 791

Internet Control Message Protocol (ICMP).
RFC 792

Transmission Control Protocol (TCP).
RFC 793

Ethernet Address Resolution Protocol (ARP).
RFC 826

Transmission of IP Datagrams Over Public Data Networks
RFC 877

Broadcasting Internet Datagrams.
RFC 919

Broadcasting Internet Datagrams in the Presence of Subnets.
RFC 922

Internet Standard Subnetting Procedure.
RFC 950

Bootstrap Protocol.
RFC 951

Using ARP to Implement Transparent Subnet Gateways.
RFC 1027

Serial Link Internet Protocol (SLIP).
RFC 1055

Routing Information Protocol (RIP).
RFC 1058

Simple Network Management Protocol (SNMP).
RFC 1155

Internet Numbers
RFC 1166

Concise MIB Definitions.
RFC 1212

Management Information Base for Network Management of TCP/IP-Based Internets: MIB II.
RFC 1213

Internet Control Message Protocol (ICMP) Router Discovery Messages.
RFC 1256

Management Information Base for Frame Relay Data Terminal Equipment (DTE), 1992.
RFC 1315

Internet Protocol Control Protocol (IPCP) for the Point-to-Point Protocol (PPP).
RFC 1332

Password Authentication Protocol (PAP)
RFC 1334

Multiprotocol Interconnect on X.25
RFC 1356

Multiprotocol Interconnect over Frame Relay Data Terminal Equipment (DTE), 1993.
RFC 1490

Definition of Managed Objects for Bridges.
RFC 1493

Internetwork Packet Exchange Protocol Control Protocol (IPXCP) for the Point-to-Point Protocol (PPP).
RFC 1552

Point-to-Point Protocol (PPP) over Integrated Services Digital Networks (ISDN).
RFC 1618

Network Address Translation (NAT)
RFC 1631

Bridge Control Protocol (BCP) for the Point-to-Point Protocol (PPP).
RFC 1638

Link Control Protocol (LCP) for the Point-to-Point Protocol (PPP).
RFC 1661

Point-to-Point Protocol (PPP) over High-level Data Link Control (HDLC).
RFC 1662

Assigned Numbers.
RFC 1700

Multilink Point-to-Point Protocol (PPP)
RFC 1717

Routing Information Protocol (RIP) Version 2
RFC 1723

Requirements for IP Version 4 Routers — regarding subnetting.
RFC 1812

Point-to-Point Protocol (PPP) Data Compression
RFC 1962

Point-to-Point Protocol (PPP) Encryption
RFC 1968

Point-to-Point Protocol (PPP) Stac Data Compression
RFC 1974

Point-to-Point Protocol (PPP) Multilink
RFC 1990

Challenged Handshake Authentication Protocol (CHAP)
RFC 1994

Triggered RIP
RFC 2091

High-level Data Link Control.
ISO/IEC 3309

Link Access Procedure, Balanced (LAPB)
ISO/IEC 7776

X.25 Packet Layer Protocol for Data Terminal Equipment (DTE).
ISO/IEC 8208

Frame Relay User-Network Interface, January 1992.
FRF.1

Frame Relay Multiprotocol Interconnect, December 1992.
FRF.3

Media Access Control (MAC) Bridges.
IEEE/ANSI 802.1D
IEEE/ANSI 802.1G

High-level Data Link Control (HDLC)
ISO 3309

Frame Relaying Bearer Service Architectural Framework and Service Description, 1990.
ANSI T1.606

Frame Relaying Bearer Service, Congestion Management Principles, 1991 Addendum 1.
ANSI T1.606

DSS1 — Signalling Specification for Frame Relay Bearer Service, 1991.
ANSI T1.617

DSS1 — Core Aspects of Frame Protocol for use with Frame Relay Bearer Service, 1991.
ANSI T1.618

Frame Relay Bearer Services, 1991.
ITU-T (CCITT) Recommendation I.233.1

Congestion Management in Frame Relaying Networks, 1991.
ITU-T (CCITT) Recommendation I.370

Basic User-Network Interface — Layer 1 Specification, 1993.
ITU-T (CCITT) Recommendation I.430

Usage of Cause and location in the Digital Subscriber Signalling System No. 1 and the Signalling System No. 7 ISDN User Part, 1993.
ITU-T (CCITT) Recommendation Q.850

ISDN User-Network Interface-Data Link Layer Specification, 1993.
ITU-T (CCITT) Recommendation Q.921

ISDN Data Link Layer Specification for Frame Mode Bearer Services, 1992.
ITU-T (CCITT) Recommendation Q.922

ISDN User-Network Interface Layer 3 Specification for Basic Call Control, 1993.
ITU-T (CCITT) Recommendation Q.931

International Numbering Plan for Public Data Networks
ITU-T (CCITT) Recommendation X.121

Interface Between Data Terminal Equipment (DTE) and Data Circuit-terminating Equipment (DTE) for Terminals Operating in the Packet Mode and Connected to Public Data Networks by Dedicated Circuits.
ITU-T (CCITT) Recommendation X.25

Part I

LAN and WAN Links and Services

Chapter 1

LAN and WAN Services in the Router

In this chapter

This chapter gives an introduction to the LAN and WAN services available in the different Intel Express 8100 Router versions, and the common facilities available of the different WAN services. WAN services include leased line, PPP, Frame Relay, ISDN and X.25 services.

The different WAN services available in the Intel Express 8100 Router are described in the following chapters.

LAN Services

Introduction

Local Area Networks (LAN) are concerned with the interconnection of distributed computer systems whose physical location is confined to a localized group of buildings. The main difference between a communication path established via a LAN and connections made via Public Data Networks/Wide Area Networks (WAN) is that connections via a LAN allow for higher transmission rates due to the short physical distances.

LAN Services concept

The following illustration indicates how LAN Services are embedded in the Intel Express 8100 Router.

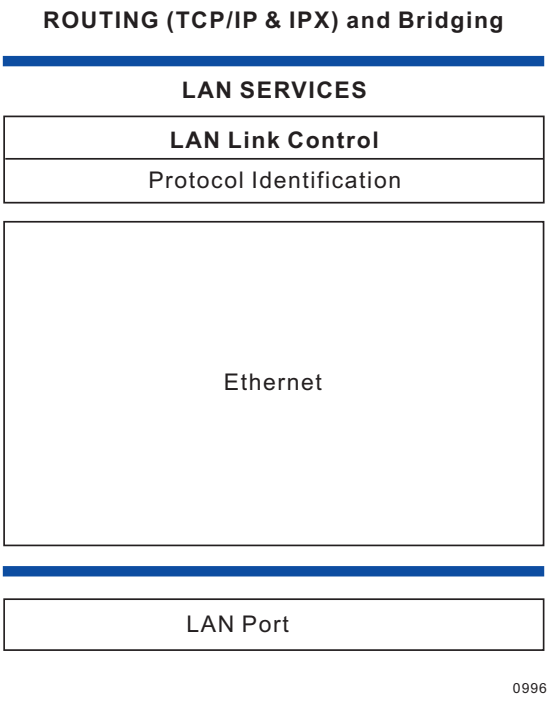

Ethernet services

Ethernet is a Local Area Network (LAN) hardware standard that is capable of linking up to 1,024 nodes (stations) in a bus network. It uses a base-band (single-channel) communication technique providing for a raw data transmission rate of either 10 Mbps for 10Base-T or 100Mbps for 100Base-T. Ethernet uses carrier-sense multiple-access/ collision-detection techniques (CSMA/CD).

WAN Services

Introduction

Wide Area Network (WAN) services consist of links to remote sites via private or Public Data Networks (PDNs).

WAN Services and Protocols Available

WAN services overview

WAN services are private or public data networks (PDNs) available to subscribers for interconnecting remote sites. The Intel Express 8100 Router supports:

- Leased lines
- Switched (dial-up) connections
- Point-to-Point Protocol (PPP)—described in "Point-to-Point Protocol (PPP)", p. 8
- Frame Relay services—described in Chapter 3 "Frame Relay Services", p. 17
- ISDN connections and services—described in Chapter 4 "ISDN Services", p. 20
- X.25 services—described in Chapter 5 "X.25 Services", p. 32
- PPP Multilinks which combine a number of PPP links into a single route between two sites—described in Chapter 6 "PPP Multilinks", p. 39
- Internet Tunnels for routing data via the internet—described in Chapter 7 "Internet Tunnels", p. 41

Number of WAN links supported

The Intel Express 8100 Router supports up to 5 WAN links (Frame Relay, ISDN or X.25).

Leased lines

A leased line is a permanent physical connection between two Local Area Networks. The costs for the use of a leased line are fixed—that is, they are independent of the amount of data transmitted. A leased line is normally used when there is a need for constant data flow between two sites.

Dial-up links

Dial-up links (also referred to as switched links or switched lines) are links that are only established when data requires transmitting over the WAN link and are similar in operation to a normal telephone connection. Dial-up WAN links require the use of a modem or other WAN connect device designed to establish this type of connection.

The number for the dial-up WAN link is configured in the modem or WAN connect device (not for ISDN and X.25 links). This is a prerequisite for V.11, V.24, V.35 and V.36 WAN connections.

General Facilities for WAN Services

Introduction

This section gives an overview and describes the WAN service features common to all WAN link types available in the Intel Express 8100 Router.

Data compression

Data compression is available for all WAN link types to improve the throughput.

PPP and Frame Relay

For PPP links (leased lines, dial-up links and ISDN links) and Frame Relay links, data compression rates of up to 4:1 can be achieved for text data.

X.25

For X.25 links, data compression rates of up to 4:1 can be achieved for text data.

Timecut mode for dial-up (switched) links

Note Timecut mode should be selected for dial-up (switched) WAN links. Operating costs can be excessive otherwise.

When a dial-up (switched) WAN link is in Timecut mode, the link is only activated when there is queued data. Timecut mode should be selected where transmission costs are dependent on the data transferred, the number of data packets sent or the duration for which the link is active. Timecut mode ensures that the WAN link is only established when necessary.

Routing protocols such as RIP and IPX send various routing packets between devices to continually assess the topology of the network. If the standard settings are used for these settings (for example RIP packets are sent every 30 seconds), timecut WAN links never get the chance to close down. Routing protocol settings used on timecut WAN links should be set in accordance with this, and there are other considerations to be made (for example Watchdog Spoofing and RIP and SAP Filtering—see Chapter 9 “Novell IPX Routing”, p. 71).

Backup links

Backup links are dial-up (switched) links that can be initialized if another link should fail or has not been established within a defined time interval. Typically, a primary WAN link is a 2 Mbps leased line connection, the backup WAN link is a 64 kbps dial-up connection—the costs of the backup link is therefore dependent on how much the link is used.

A WAN link in Backup mode is only used if the primary link should fail, or if a connection to it was established from a remote source. Backup links can only be closed down by the router that started the link.

Timecut backup links considerations

Timecut mode for backup links must only be used if the backup link runs in parallel to the primary WAN link, that is the main WAN link and the backup WAN link must run between the same routers as shown in the figure below:

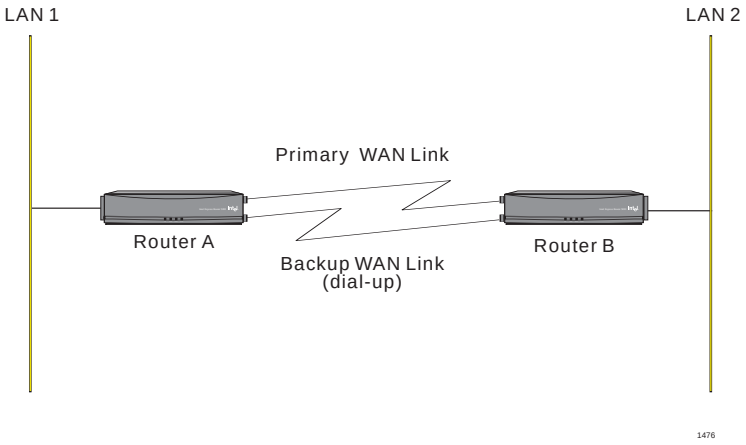

Timecut mode MUST NOT be selected for backup links that run via a different router than the main WAN link as shown below. Problems otherwise occur when using the backup WAN Link.

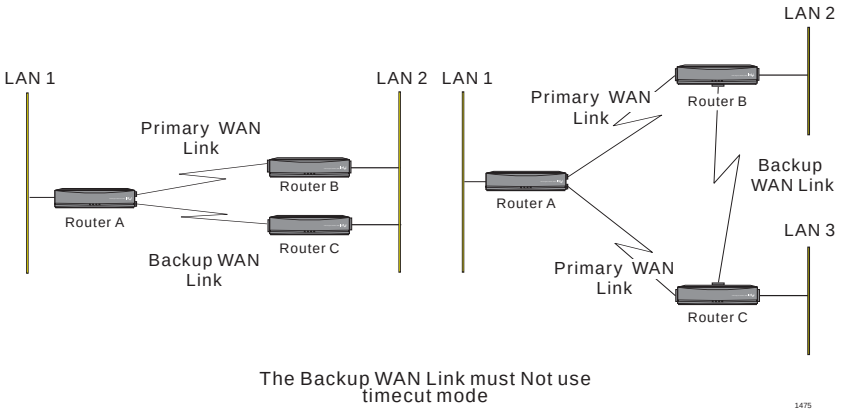

Multiple backup links

For ISDN and X.25 services, it is possible to establish multiple backup links for increased internetwork reliability. This can be a particularly useful feature in X.25 networks where a number of Switched (dial-up) links can be available.

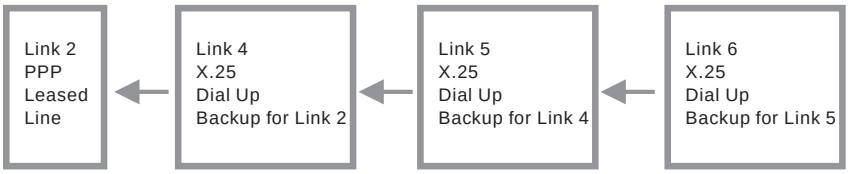

In the above example, link 4 is used if the primary link (link 2) should fail, link 5 is used if link 4 should fail and link 6 is used if link 5 should fail.

Timer Profiles

The Intel Express 8100 Router implements Timer Profiles which offer extensive facilities to restrict WAN link activity according to a user-defined time profile. This may be done for security reasons or to reduce the operating costs of WAN links.

Up to 16 user-defined timer profiles can be defined, any one of which can be selected for use with a WAN link during configuration of the link. Each profile allows you to define access rights on a weekly basis, with a half-hour resolution. For each link, access can be blocked:

- for outgoing access from the router to the WAN link
- for both outgoing and incoming access (WAN link disabled outside the times defined in the timer profile)

For example, for one link using a timer profile, outgoing access may be block, while for another WAN link using the same timer profile, access may be block for all data traffic.

Timer profiles consideration

Warning When blocking both outgoing and incoming access, the timer profile must be the same in both the routers over the WAN link.

This is particularly important for dial-up (switched) WAN links where operating costs are dependent on the use of the link.

If the timer profiles are not the same and only one of the routers is denying incoming access at some time, the other router may continually attempt to establish the WAN link. This can lead to the following problems for the calling router:

- for dial-up WAN links, excessive operating costs for the WAN link
- error messages being logged
- the link being disabled (faulted)

Example use of a timer profile

For example, a timer profile could be set up to deny both incoming and outgoing access on a WAN link outside normal office hours of 7:30 to 17:30, and to deny all access at weekends.

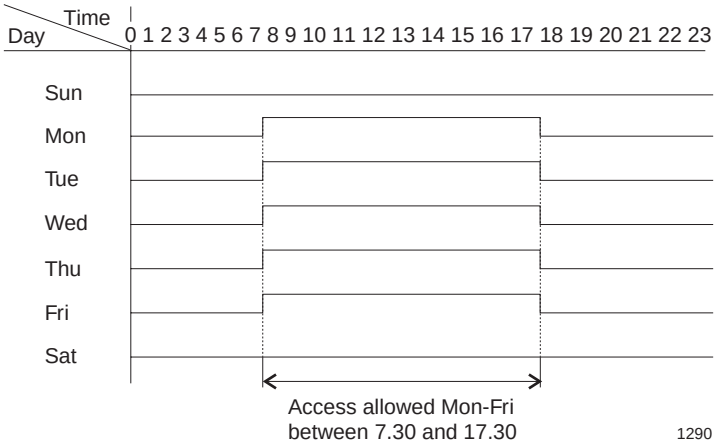

Daily Activity Limit

The Intel Express 8100 Router has a Daily Activity Limit function which can be used to control the use, and therefore the cost of operation, of dial-up WAN Links. The alarm generates an SNMP Trap and can also be set up to close and lock the associated WAN link when it has been in the Up (for both call directions) state for the configured time-period within a day (midnight to midnight). That is, the activity counter is incremented for both incoming and outgoing calls on the links. If the link has been locked by the Daily Activity Limit, it stays locked until it is manually reset from Intel Device View.

When a dial-up link has been locked by the Daily Activity Limit, the router does not permit outgoing calls on the link until the link has been unlocked again from Intel Device View. However, the link accepts incoming calls, thereby allowing a remote connection via the link to unlock the link.

Chapter 2

Leased Lines Links

In this chapter

The WAN connect protocol used for communicating over leased line WAN connections by the Intel Express 8100 Router is the Point-to-Point Protocol (PPP). When communicating over leased line connections, PPP runs on top of the High-level Data Link Control (HDLC) protocol.

Note PPP is also used for communications over ISDN links (PPP over ISDN).

This chapter gives an introduction to the Point-to-Point Protocol (PPP) and describes the features of the protocol offered by the Intel Express 8100 Router.

Point-to-Point Protocol (PPP)

Introduction to PPP

The Point-to-Point Protocol (PPP) was designed to enable simultaneous transfer of network-layer protocols across a point-to-point link. Its main function is to establish a synchronous link connection between routers from any manufacturer. PPP functions include:

- encapsulating multi-protocol datagrams
- establishing, configuring and testing the data-link connection using a Link Control Protocol (LCP)
- establishing and configuring bridging and various network-layer communications across the link using Network Control Protocols (NCPs). This is handled by the Internet Protocol Control Protocol (IPCP) for IP, and the Internetwork Packet Exchange Control Protocol (IPXCP) for IPX and the Bridging Control Protocol (BCP) for bridging services
- authenticating peers using the Challenge Handshake Authentication Protocol (CHAP) and Password Authentication Protocol (PAP) to ensure that communications are between authorized devices
- encrypting and decrypting data transmitted via the link to offer data security
- compression/decompression of data to improve the throughput of data

PPP provides transport services for data packet delivery with low overheads and high throughput. Frame checking at the link-level offers error detection, but error recovery is taken care of by higher-layer network protocols.

PPP protocols

The following diagram shows how the various PPP protocols are embedded:

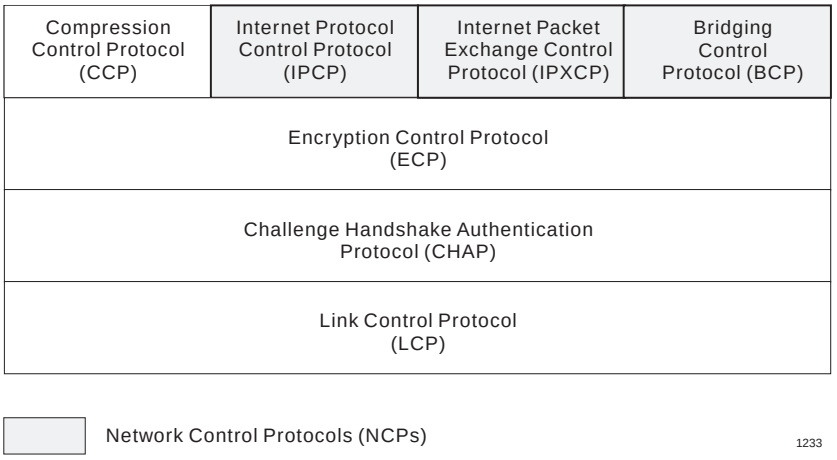
Compression Control Protocol (CCP)
Internet Protocol Control Protocol (IPCP)
Internet Packet Exchange Control Protocol (IPXCP)
Bridging Control Protocol (BCP)
Encryption Control Protocol (ECP)
Challenge Handshake Authentication Protocol (CHAP)
Link Control Protocol (LCP)
Network Control Protocols (NCPs)
1233

Link Control Protocol (LCP)

LCP is the lowest layer in the PPP stack and runs on top of the High-level Data Link Control protocol (HDLC) or Integrated Services Digital Networks (ISDN). When the HDLC or ISDN protocol has established the link, LCP runs on top of it.

Operation of PPP

The following protocols form the basis of PPP:

Encapsulation

Network-layer protocols are encapsulated for transmission over WAN links.

Link Control Protocol (LCP)

The LCP takes care of the link connection.

Challenge Handshake Authentication Protocol (CHAP)

The Challenge Handshake Authentication Protocol (CHAP) is available to offer security (encrypted password protection) against unauthorized access to a PPP link.

Password Authentication Protocol (PAP)

The Password Authentication Protocol (PAP) is available as an alternative to CHAP offering security (password protection) against unauthorized access to a PPP link.

Encryption Control Protocol (ECP)

Data encryption is available when communicating over PPP links. The Encryption Control Protocol (ECP) negotiates and manages data encryption between the devices over the link.

Note Data encryption is only available in certain models of the Intel Express 8100 Router. Data encryption is not allowed in some countries by law.

Compression Control Protocol (CCP)

Data compression is available when communicating over PPP links. The Compression Control protocol (CCP) negotiates and manages data compression between devices over the link.

Network Control Protocols (NCPs)

Each Network Control Protocol negotiates and manages a specific network-layer/bridging protocol.

PPP Call Back (PPP ISDN links)

PPP Call-Back is available for use over PPP ISDN links where the costs of operation of the link can be transferred a specific location.

For example, if someone working from home needs an occasional connection to the office, the costs of operation for the connection can be transferred to the office by using Call-Back. When the home connection needs to communicate with the office, they call in to the office Intel Express 8100 Router with a request for call back and then closes the connection.

Note The router will then return the call. The router must support Call-Back or the link will not work.

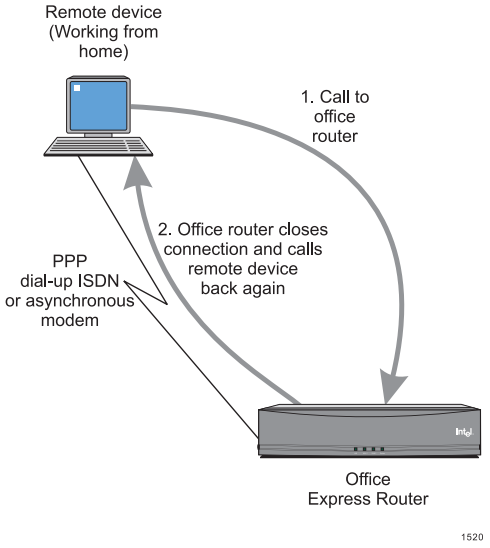

In some cases, Call-Back is used to verify that incoming call is from the authorized address set up for the link.

The Call-Back facility can be used to transfer the cost of operation to either the local site (this router) or to the remote site.

Multilink PPP

The Intel Express 8100 Router implements PPP Multilink facilities which allow multiple PPP links between two sites to be used as a single route. A PPP Multilink allows links to added dynamically (for bandwidth-on-demand) or statically and has almost the combined bandwidth of the individual links.

PPP Multilinks are described in Chapter 6 “PPP Multilinks”, p. 39.

Data Compression

Data compression

Data compression is available when communicating over PPP links, to improve the throughput.

Both routers over the PPP link must support and be configured for PPP data compression for data compression to be used.

Compression algorithm used

The compression algorithm implemented in the Intel Express 8100 Router is based on the Stac* algorithm developed by Stac Electronics Inc. Typical compression rates of 4:1 are achieved for text data.

Requirements for use

The devices at both end of the PPP link must implement the Compression Control Protocol (CCP) and use the same compression algorithm.

Data compression is negotiated by the CCP whenever the link is established. If the device over the PPP link does not support the CCP or the same compression algorithm, compression cannot be used over the link.

Link speeds

Data compression can be used on combined link speeds of up to 256 kbps (for example compression can be used on one PPP link of 256 kbps or on two PPP links of 128 kbps). The delays involved with compressing the data make it impractical to use at higher link speeds.

Data Encryption

Introduction

The Intel Express 8100 Router offers encryption/decryption of data being transmitted over PPP and Frame Relay links. This offers security in case of interception by an unauthorized source.

Note Data encryption is only available in certain models of the Intel Express 8100 Router which are not available in some countries.

Encryption algorithm used

The data encryption algorithm implemented in the Intel Express 8100 Router is Blowfish with a variable length encryption key (up to 144 bits) with 16 rounds (encryption steps). The algorithm is used in Cipher Block Chaining (CBC) mode

which means that the algorithm is used across the entire data stream including the packet header containing the address and protocol identification, and not only on a fixed block (packet) size.

Reference for the encryption algorithm

The Blowfish encryption algorithm is described in:

Bruce Schneier
Applied Cryptography (John Wiley & Sons)

Requirements for use

The devices at both end of the PPP link must implement the Encryption Control Protocol (ECP) and use the same encryption algorithm.

Data encryption is negotiated by the ECP whenever the link is established. If the device over the PPP link does not support the ECP or the same encryption algorithm, the link is disconnected and a message is entered in the System Log for the router—data communications are not allowed on a PPP link intended for secure communications.

Link speeds

Encryption can be used on all link speeds and can also be used in conjunction with compression. The algorithm can encrypt at around 1.3 Mbps, which may cause delays on combined link speeds above this (for example on a 2.0 Mbps links).

Encryption and compression

Data encryption can be used together with data compression (see "Data Compression", p. 11) over a PPP link. Data is first compressed then encrypted. When encryption is used in connection with data compression over a PPP link, the restrictions on link speeds for data compression apply.

Peer Authentication using the Challenge Handshake Authentication Protocol (CHAP)

Introduction to CHAP

The Challenge Handshake Authentication Protocol (CHAP) can be used to provide link security against unauthorized access. CHAP uses password encryption where passwords can be global (used for all PPP links) or selected from a password pool. Separate passwords can be used for incoming and outgoing calls on a link.

Passwords

CHAP uses password encryption to authenticate peers; separate passwords can be used for incoming and outgoing calls. The passwords are used to encrypt random text files which are transmitted over the PPP link (see 'Challenge handshake authentication procedure' following); passwords are therefore never transmitted directly over a PPP link, and cannot be intercepted and used by unauthorized sources.

Use of passwords

For the correct operation of PPP links using CHAP, the passwords must be configured as follows:

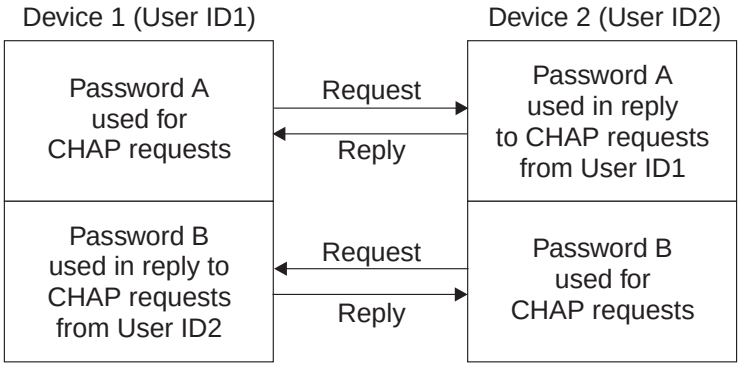
Device 1 (User ID1)
Device 2 (User ID2)
Password A used for CHAP requests
Request
Reply
Password A used in reply to CHAP requests from User ID1
Password B used in reply to CHAP requests from User ID2
Request
Reply
Password B used for CHAP requests
1338

If Device 1 has Password A configured for CHAP requests on the PPP link, Device 2 must reply to the request using Password A. That is, Password A must be defined in Device 2 for CHAP replies to Device 1.

Similarly, if Device 2 has Password B configured for CHAP requests on the PPP link, Device 1 must reply to the request using Password B. That is, Password B must be defined in Device 1 for CHAP replies to Device 2.

Global (all PPP links) and Local replies to CHAP requests

A device always tries to reply to a request for authentication from a peer, using the password defined for the User ID of the peer.

In the Intel Express 8100 Router, a list of passwords for User IDs can be defined for both the router (global) and for individual PPP links. The passwords in the global list can be used in reply to CHAP requests from peers on any of the PPP links to the router. The passwords in the local list can only be used for peers on the PPP link for which they are defined. The router first tries to find the User ID of the peer requesting authentication in the local password list and then in the global password list.

In some devices, a common User ID with different passwords are used for communications over a PPP link. In these cases, local User IDs and password should be defined. Otherwise, global User IDs and passwords can generally be used.

CHAP requests from the router

When CHAP requests are enabled, the Intel Express 8100 Router authenticates the peer over the PPP link whenever the link is established. For on-demand (switched) PPP links including ISDN links, the router authenticates the peer whenever the link is brought up.

Challenge handshake authentication procedure

The procedure used to authenticate a peer over a PPP link is as follows:

1 The device which is authenticating the peer (Device 1) generates a random text file (random characters and a random length), and transmits it to the device to be authenticated (Device 2) over the PPP link. Device 1 also stores a copy of the random text file which is encrypted using the password defined for CHAP requests on the link (Password 1).

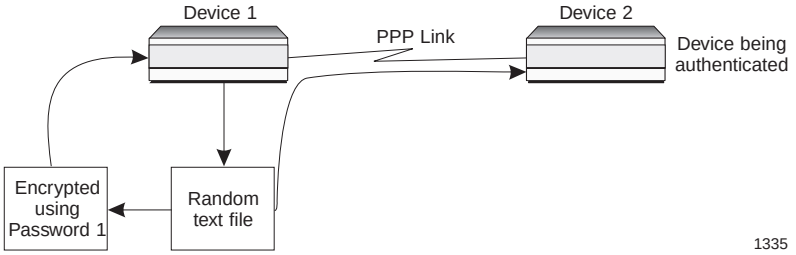

2 The device receiving the random text file (Device 2) encrypts it using the password (Password 2) defined for requests from Device 1, then transmits the encrypted text file back to Device 1.

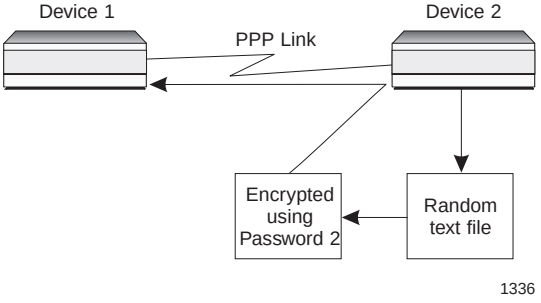

3 Device 1 compares the locally encrypted text file with the text file encrypted by Device 2. If the text files are the same (Password 1 = Password 2) Device 2 is authenticated. Otherwise, Device 2 is not authenticated and data communications with Device 2 are not allowed.

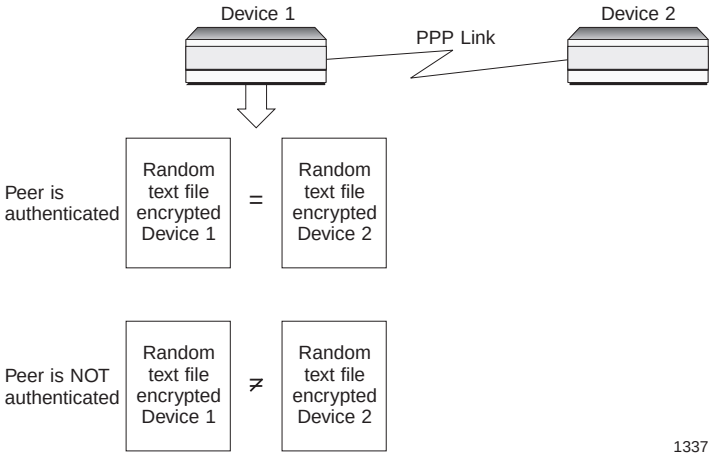

Peer Authentication using the Password Authentication Protocol (PAP)

Introduction to PAP

The Password Authentication Protocol (PAP) can be used as an alternative to CHAP to provide link security against unauthorized access. PAP uses simple password protection against unauthorized access.

PAP versus CHAP

CHAP uses password encryption to authenticate peers and the passwords are never transmitted directly over the PPP link, and therefore cannot be intercepted and used by unauthorized sources. PAP uses simple password protection where the password is transmitted directly over the link—PAP is therefore not resilient to link monitoring.

Note For the best security, you should use CHAP rather than PAP for link access protection wherever possible.

PAP can however be used when it is the only authentication protocol supported by the remote device.

Replies to PAP requests

A device (for example the Intel Express 8100 Router) always tries to reply to a request for authentication from a peer, using the password defined for the User ID of the peer.

PAP requests from the router

When PAP requests are enabled, the Intel Express 8100 Router authenticates the peer over the PPP link whenever the link is established. For on-demand (switched) PPP links including ISDN links, the router authenticates the peer whenever the link is brought up.

If the device over the PPP link fails to authenticate itself, communications over the link are not allowed.

Chapter 3

Frame Relay Services

Frame Relay in the Intel Express 8100 Router

Frame Relay is available in the Express 8100 Router with Frame Relay. The WAN port must be configured as a Frame Relay port before it can be used for Frame Relay services.

As Frame Relay uses Permanent Virtual Circuits, a number of Frame Relay links can be made via the same Frame Relay port. The Intel Express 8100 Router supports up to 5 Frame Relay links.

This chapter gives an introduction to Frame Relay services in the Intel Express 8100 Router.

Frame Relay

Frame Relay is an interface specification that provides a signalling and data transfer mechanism between endpoints (routers) and the network (switches). Data to be transmitted are encapsulated according to the Multiprotocol Encapsulation Implementation Agreements of the Frame Relay Forum. It provides efficient network services and accommodates burst-intensive applications over wide area LAN interconnections at rates up to 2.048 Mbps. Frame Relay provides bandwidth on demand and allows multiple simultaneous data sessions (logical channels) across a physical Frame Relay port.

Frame Relay can be very cost-effective when used to interconnect several sites. Frame Relay allows several logical links to different sites across one physical Frame Relay port.

Features

Features of the Frame Relay protocol include:

- Provides a connection-oriented service
- Maintains data sequence
- Connects end-systems at data link layer
- uses variable length packets
- no data link control, no retransmissions
- makes use of LAPF core functions
- standard multiprotocol encapsulation
- reduced costs for internetworking in large networks

- increased interoperability as it complies with international standards

Addressing

Each data packet contains a circuit number, also called a Data Link Connection Identifier (DLCI), showing which logical channel that the information belongs to. Frame Relay packets are routed to their destination on the basis of the circuit numbers contained in the packet.

Note When using Frame Relay, ensure that your network provider allocates you the necessary DLCI values. DLCI values assigned during Frame Relay connection, range from 16 up to 991.

Encapsulation

Multiprotocol encapsulation according to FRF.3, specifying the protocol contained in the current Information field (I-field). This encapsulation is also specified in RFC-1490 “Multiprotocol Interconnect over Frame Relay Data Terminal Equipment”, and is compatible with routers from other manufacturers.

PVC management

Status information between Frame Relay network switches and user devices (for example routers and bridges) is provided by signalling protocols defined by the Frame Relay standards committees. These signalling protocols perform the following tasks:

- providing a ‘heartbeat’ or ‘keep-alive’ message exchange to ensure that links are available
- informing about added and removed PVCs
- providing status of existing PVCs

Frame Relay elements

Frame Relay services may be broken down into three distinct elements:

Frame Relay access equipment

Frame Relay access equipment is customer premises equipment (CPE) which sends information across the network. This may be, for example, a router, a bridge, a host computer, and so on.

Frame Relay switching equipment

Frame Relay switching equipment is responsible for the transportation of information across a WAN. This may be, for example, T1 (1.5 Mbps) and E1 (2 Mbps) packet switches, and so on.

Public Frame Relay services

Public Frame Relay services see the deployment of Frame Relay switches via the Public Data Network (PDN). If you do not make use of public services (subscribed), private networks may be established. In this case, the private network must have Frame Relay switching equipment installed.

Note Frame Relay can only be used on high-quality transmission lines dedicated to Frame Relay.

Data compression

Data compression is available when communicating over Frame Relay links, to improve the throughput on link speeds of up to 256 kbps. Typical compression rates of 4:1 can be achieved for text data.

Both routers over the Frame Relay link must support and be configured for Frame Relay data compression before data compression is used.

Intel Express 8100 Routers use the Stac algorithm developed by Stac Electronics Inc. Typical compression rates of 4:1 are achieved for text data.

Data encryption

The Intel Express 8100 Router offers encryption/decryption of data being transmitted over Frame Relay links. This offers security in case of interception by an unauthorized source.

Note Data encryption is only available in certain models of the Intel Express 8100 Router which are not available in some countries.

Data encryption over Frame Relay links uses the same encryption algorithm and implementation as for PPP links, see "Data Encryption", p. 11.

Note Unlike PPP links, Frame Relay links are unable to detect a mismatch in the encryption keys in the routers over the link. If the encryption keys are wrong, data will not be un-encrypted and will contain errors (illegal packets).

Advantages of Frame Relay over leased line links

Frame Relay uses standardized encapsulation of protocols and it is therefore possible to connect to third-party routers (routers from other manufacturers). When using Frame Relay, virtual circuits are established allowing for any-to-any connectivity. Both of the mentioned features are not available for a leased line.

Chapter 4

ISDN Services

Integrated Services Digital Network (ISDN) in the Intel Express 8100 Router

ISDN in the Intel Express 8100 Router

Basic Rate Access (BRA) ISDN is supported by the following routers:

Express 8100 Router with an ISDN S/T port

which establishes a Basic Rate Interface (BRI) via an ISDN S/T inteface to an NT-1 termination of the ISDN line

Express 8100 Router with an ISDN U port

which establishes a Basic Rate Interface (BRI) via an ISDN U interface directly to the incomming ISDN line.

Point-to-Point Protocol (PPP)

The link protocols used over ISDN links by the Intel Express 8100 Router is the Point-to-Point Protocol (PPP). PPP is also be used for communicating over leased line and dial-up (switched) links.

An overview of PPP and the facilities available is given in Chapter 2 “Leased Lines Links”, p. 8.

ISDN links

For the ISDN Intel Express 8100 Routers, it is possible to configure more ISDN links (router links) than the number of B-channels available. In this case, ISDN B-channels can be allocated to ISDN links as follows:

Reserved ISDN B-Channels

ISDN links in the router can be configured to reserve an ISDN B-channel. In this case, the ISDN link reserving the B-channel always has a B-channel available. The B-channel reserved by the ISDN link cannot be allocated to another ISDN link when not in use.

Dynamically allocated ISDN B-Channels

For ISDN links which do not have a reserved ISDN B-channel, a simple link pool system is used. ISDN B-channels which are not reserved by ISDN links are obtained and returned to the B-channel pool as required by the ISDN links.

ISDN Concept

What is ISDN?

ISDN is a digital telecommunications network that allows you to send all kinds of information, including speech and data, in much the same way as making a telephone call. The ISDN is digital from end-to-end which improves the data communications quality and restricted information rate resulting from the use of analog links together with modems.

CCITT has defined an ISDN as a network, in general evolving from a telephony Integrated Digital Network (IDN), that provides end-to-end digital connectivity to support a wide range of services, including voice and non-voice services, to which you have access by a limited set of standard multi-purpose user-network interfaces.

What ISDN offers

ISDN offers multiple 64 kbps or 56 kbps connections on demand with a low error rate. This can be used in place of leased line WAN links or as a backup to leased line WAN links. A main advantage of ISDN is that a modem is not required—the router can be connected directly to the ISDN.

Channel types

There are two types of channels used in ISDNs:

- B-channels
  64 kbps or 56 kbps full duplex, synchronous channels for data and voice transfer
- D-channels

  16/64 kbps signalling channels

Connection types

There are two types of connectivity:

Basic Rate Access (2B+D)

This type of connection provides two independent data (B) channels (56 kbps or 64 kbps), and a 16 kbits signalling (D) channel which is provided for communication between the router and the ISDN exchange. It is used to set up and clear calls for both information channels. This arrangement is called 2B+D and is equivalent to two independent telephone lines each with a 64 or 56 kbps capacity. Up to eight terminals can be connected to the line, although only two calls can be in progress at any one time (each using a B-channel).

Primary Rate Access (30B+D)

This type of connection is used for connecting Private Branch Exchanges (PBX) to the ISDN. In most countries, Primary Rate Access has 30 B-channels (56 kbps or 64 kbps) and one D-channel (64 kbps). In the USA and Japan a 23B+D arrangement is used.

Note The Intel Express 8100 Routers offer Basic Rate Access (BRA) only.

ISDN Services

Three types of service

ISDN services are divided into three categories:

- Bearer Services
- Tele Services
- Supplementary Services

Bearer Services

These are the basic transport services which transfer information between the terminal/net interface. That is a connection must be set up between two routers before communication can take place. There are two types of Bearer services:

Circuit switched

These are the simplest types of Bearer services. Once a connection is established, data is transferred over the link (B-channel) at a rate of 64 kbps (or 56 kbits for some ISDNs).

Packet switched

If the data rate is variable, then it may not be cost-effective to pay for an end-to-end channel for the duration of the call. In this case, data is sent in packets and the cost is based, partly, on the volume of traffic rather than the duration of the call.

Tele Services

These define the way in which certain standardized applications use ISDN. The following services are defined by CCITT:

Telephony

Provides two-way, real-time speech conversation. There are two types:

- 3.1 kHz
- 7 kHz, which provides speech and sound transmission of a much higher quality than 3.1 kHz telephony

Teletex

Enables subscribers to exchange character-coded documents automatically. It is much faster than Telex and has a much larger character set.

Telefax 4

Provides high-speed, high-quality facsimile transmission.

Mixed Mode

This is a combination of Teletex and Telefax 4.

Videotex

Provides retrieval and mailbox functions for text and graphic information.

Telex

Provides interactive text communication. It is the same as the existing Telex service. The advantage of using it via ISDN is that it can be provided over an interface shared with other services.

Message Handling System (MHS)

Allows you to send and receive electronic mail via a “post box”.

Supplementary Services

A Supplementary Service modifies or supplements a basic service. It may provide additional control functionality or provide information about an ISDN call. Some of the Supplementary Services available are:

Multiple Subscriber Number (MSN)

Allows more than one number to be assigned to a line. Typically, a block of ten numbers is assigned with the last number used to select the station.

Sub-addressing (SUB)

Allows you to specify extra addressing information in addition to the ISDN number.

Direct Dialing In (DDI)

Allows an ISDN station to call on an ISPBX directly without the operator.

User-to-User Signalling (UUS)

Allows you to send or receive a limited amount of information to or from another station over the D-channel in association with an existing call.

Calling Line Identification Presentation (CLIP)

The called party can choose to be presented with the calling party’s ISDN number.

Calling Line Identification Restriction (CLIR)

The calling party can suppress the presentation of his ISDN number to the called party.

Connected Line Identification Presentation (COLP)

The calling party can be presented with the ISDN number of the called party. It tells the calling party who actually answers the call (for example if call forwarding is in operation).

Connected Line Identification Restriction (COLR)

The called party can suppress the presentation of his ISDN number to the calling party.

Closed User Group (CUG)

Allows a group of users to appear to be in their own private network.

Terminal Portability (TP)

Allows a terminal to be unplugged from one socket on a Basic Rate Access connection (S bus) and plugged into another socket on the same S bus without the call being cleared.

Call Waiting (CW)

Gives notice of an incoming call even if all channels are occupied.

Call Hold (HOLD)
Allows you to interrupt a call and return to it later.

Line Hunting (LH)
Allows incoming calls to a specific number to be distributed over a group of lines.

Call Forwarding (CF)
Causes an incoming call to be sent to another number.

Call Transfer (CT)
Allows either the calling or called party to transfer the call to a third party.

Add on Conference Call (CONF)
Allows additional subscribers to be added to a call to form a conference.

Three-Party Service (3PTY)
Allows you to switch between two calls.

Advice of Charge (AOC)
Gives call charge information.

ISDN Numbering and Addressing

Introduction to ISDN numbering

Like the Public Switched Telephone Network (PSTN), Integrated Services Digital Networks (ISDNs) uses Directory Numbers (DNs) to identify subscribers. The International Telegraph and Telephone Consultative Committee (CCITT) a numbering plan for ISDN (Recommendation E164) which specifies how numbers should be constructed within telephone networks. E164 includes the existing plan E163 for the PSTN—both ISDN and PSTN numbers can therefore be allocated out of the same range of numbers.

Each ISDN directory number specifies a point of connection to the network (an ISDN line). Each line can, however, support more than one terminal—see "BRA and PRA Numbering", p. 26

E164 format

The number consists of a maximum 15 digits and is divided up in the following way:

Country Code (CC)
This consists of 1-3 digits.

National Significant Number (NSN)
This consists of between 12-14 digits. This is sub-divided further into a:

- National Destination Code (NDC)
- Subscriber Number (SN)
  In Denmark, for example, the NSN is 8 digits long; NDC = 2 digits and the SN = 6 digits.

Sub Address (SA)
This consists of between 1-40 digits.

Addressing

An incoming call can contain the following address fields:

Calling Party Number
Local Address of the `Calling` router (address corresponds to the Remote Address of the `Called` router).

Calling Party Sub-address
Local Sub-address of the `Calling` router (address corresponds to the Remote Sub-address of the `Called` router).

Called Party Number
Local Address of the `Called` router (address corresponds to the Remote Address of the `Calling` router).

Called Party Sub-address
Local Sub-address of the `Called` router (address corresponds to the Remote Sub-address of the `Calling` router).

Some of the address fields may be absent depending on the network and supplementary services available.
The router accepts the call when both the destination and source address checks succeed.

Address checking

To ensure a one-to-one relationship between end devices, a unique ISDN number is assigned to each link across the ISDN. Incoming calls are only accepted if the local and remote addresses and sub-addresses match as shown below:

Addresses must match for a call to be accepted	
Calling Device	**Receiving Device**
Local Address	Remote Address
Remote Address	Local Address
Remote sub-address	Local Sub-address

Note ISDN addresses can be changed by intermediate ISDN switches in the ISDN network.

Some of the addresses may not be used by an ISDN variant particular to some countries (for example the sub-address), in which case the check of that address is omitted.

Destination address and sub-address check

The router only accepts an incoming call when the ISDN `Remote Address` matches the `Local Address`, and the `Remote Sub-address` (if used) matches the `Local Sub-address` of an ISDN link in the router.

If a `Local Address` is not defined for an ISDN link, all incoming addresses are accepted by the link (the address check is not performed). Similarly, if a `Local Sub-address` is not defined for an ISDN link, all incoming sub-addresses are accepted by the link (the sub-address check is not performed).

If the network does not provide the sub-address service, then a `Local Sub-address` must not be defined (left empty).

Source address check

Source address checking adds security to the WAN link—the router only accepts an incoming call when the source address matches the `Remote Address` configured for the ISDN link.

ISDN switches may change address prefixes. To support all ISDN implementations, the address check is made from right to left using the number of digits defined for the `Remote Address Check` parameter.

There are occasions when the source address and/or sub-address is missing from incoming calls (some ISDN networks do not provide presentation of the source address in incoming calls). In these cases, the router omits the check of the source address and accepts the calls—there is no security check on these calls.

BRA and PRA Numbering

More than one terminal supported by ISDN line

The ISDN Directory Number (DN) specifies a point of connection to the network (an ISDN line). Each ISDN line, however, can support more than one terminal. This is done in two different ways:

BRA (Basic Rate Access)
Each ISDN line can support more than one terminal directly.

PRA (Primary Rate Access)
Each ISDN line can support more than one terminal through an ISPBX (Integrated Services Private Branch Exchange) or multiplexer.

There are therefore a number of mechanisms specified to allow an incoming call to be directed to a particular terminal on a given line.

Either B-channel (if both are free) for incoming call (BRA)

Normally a BRI has a single DN and the exchange can use either B-channel (if both are free) to deliver an incoming call to a terminal.

Any free B-channel for incoming call (PRA)

Normally a PRI has a single DN and the exchange can use any free B-channel for an incoming call.

MSN (Multiple Subscriber Number)

Different Directory Number (DN) for each terminal

Each line is allocated a block of DNs, typically 10, and when there is an incoming call, the called number is presented to all the terminals on a passive bus.

Each terminal is programmed to respond to a different DN and thus the correct terminal is able to answer the call.

More than one call at same time

If the attached device is capable of handling more than one call at the same time, for example an Intel Express 8100 Router with multiple ISDN links, then it can be programmed to respond to more than one number.

Calling Line Identification

The Calling Line Identification supplementary service is a country-dependent default service for EuroISDN. A `Calling` terminal can present its DN to the network when setting up a call. This is checked by the network and, if valid, passed on to the `Called` party.

SUB (Sub-addressing)

One or more identifiers for each terminal

This is the most flexible form of terminal identification. Each terminal on a line is given one or more identifiers, each of which is a unique string of characters of your choice.

Each of the terminals on a passive bus has its own sub-address, and all of the terminals share the same ISDN number. An Intel Express 8100 Router with multiple ISDN links must have a separate sub-address for each ISDN link.

String supplied by caller

In order to specify a particular terminal or an ISDN link in the Intel Express 8100 Router, the caller supplies the appropriate character string to the network when the call is made.

The network passes this string unchanged (subject to any restrictions on length and character set) to the called line.

DDI (Direct Dialing In)

ISPBX extension numbers

This is a facility whereby the last few (typically 3) digits of the Directory Number (DN) are used to specify which extension of an ISPBX is being called.

The BRA/PRA (or group of BRA/PRAs) to which the ISPBX is connected is allocated a block of numbers.

When there is an incoming call, the last few digits of the number are presented to the ISPBX to allow it to select the correct extension.

Example

A customer could be allocated all numbers in the range 678000 to 678999.

A call to any of these numbers would be routed to the same BRA/PRA, the ISPBX being responsible for directing the call to the appropriate extension, these being numbered from 8000 to 8999.

Calling Line Identification

The Calling Line Identification supplementary service is a default service for EuroISDN. For outgoing calls, the ISPBX can provide the DDI number of the originating extension for presentation to the called subscriber.

Permanent ISDN

Permanent ISDN connections

In some countries, Permanent links are available over ISDNs which offer a static association between two sites without the need for the signalling associated with normal ISDN links. Permanent link are allocated by the ISDN provider and are similar to leased line WAN links as they are always available for data transfers.

Signalling

No signalling or dialing is required for Permanent ISDN connections. Permanent ISDN links operate on Layer-1 only—Layer-2 and Layer-3 are not used by Permanent ISDN.

Permanent ISDN link types available in the router

The Permanent ISDN link types available in the Intel Express 8100 Routers with a Basic Rate Interface (BRI) are:

- Permanent 64 kbps or 56 kbps
  which offers two 64 kbps or 56 kbps Permanent ISDN links which can be to different sites.
- Permanent 128 kbps or 112 kbps

  which offers a single 128 kbps or 112 kbps link to another site.

Permanent ISDN connections using the D-channel are not supported.
These modes of operation for the ISDN Basic Rate Interface are selected on the Advanced screen for the ISDN port.

ISDN Standards

A common standard

It should be possible to buy a piece of equipment anywhere in the world, plug it in anywhere in the world and communicate with anywhere in the world. This is the main goal of ISDN. However, at the present time, a world-wide ISDN standard does not exist. Each country has its own version of an ISDN standard. Progress is however being made, especially in Europe.

ETSI

The European Telecommunications Standards Institute (ETSI) was set up to produce European Telecommunications Standards (ETS). Some of these standards have been designated as NETs (Normes Européen de Télécommunications) which specify the connection requirements for terminal equipment. The most important are NET 3 and NET 5. These standards are for the type approval of equipment for connection to ISDN BRA and PRA respectively.

EuroISDN and National ISDN-1

In 1989, European network operators signed a Memorandum Of Understanding (MOU) to enable the standardization of services and user-network interfaces, and international connectivity.

A variety of national standards currently exist for ISDN in Europe. These are gradually being replaced by ETSI standards to provide a European-wide service. The situation is similar in the USA, a similar state of affairs exists resulting in the National ISDN-1 standard.

ISDN protocol variants supported

A number of country-specific variants of ISDN are currently in use around the world. The ISDN variants supported by the Intel Express 8100 Router are:

- Euro ISDN (ETSI)
  EuroISDN is used in many countries throughout Europe and Asia Pacific and in Australia.
- National ISDN-1 (USA)
- National ISDN-2 (USA)
- AT&T ISDN (USA)
- Nortel DMS-100 (USA)
- NTT (Japan)
- KDD (Japan)

Service Profile IDs (SPIDs)

Service Provider IDs (SPIDs) are used by USA ISDN variants only and are used to identify to the switch the ISDN services provisioned for the terminal. The SPID is registered at the switch during link establishment.

Network providers which supply NI1 (National ISDN-1) or AT&T ISDN in the USA require the use of SPIDs for identification of the services provided. For Northern Telecom DMS-100 ISDN switches, a SPID is required for each B-channel of the BRI. That is, if both B-channels of the BRI are to be used, two SPIDs are required.

ISDN Network Interface Reference Configurations

Purpose

Reference configurations are used to describe various possible physical interfaces to the ISDN.

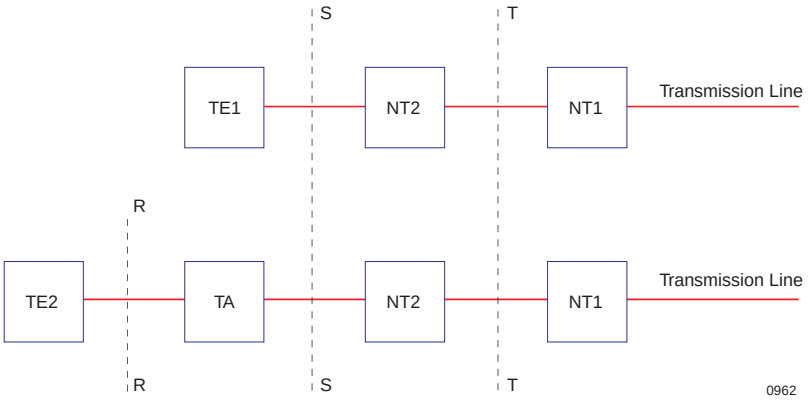

The boxes represent functional groups, each having a set of functions that may be needed to connect to the ISDN. Reference points (R, S and T) are connections between functional groups.

The Intel Express 8100 Router with an ISDN S/T port contains the following physical interfaces:

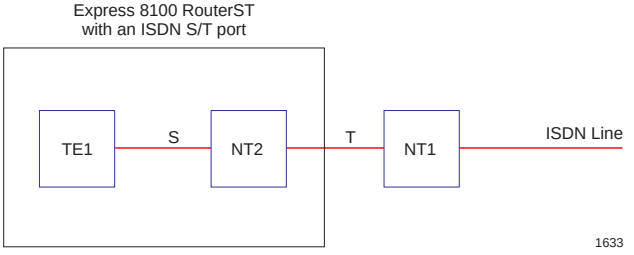

The Intel Express 8100 Router with an ISDN U portcontains the following physical interfaces:

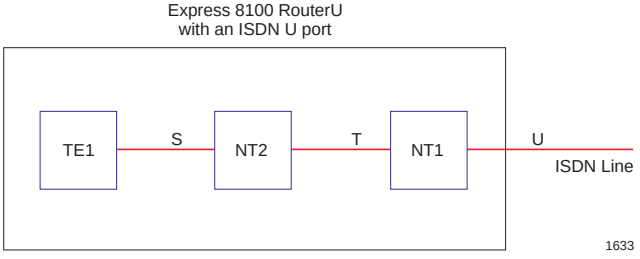

Functional Groups

The functional groups are:

- NT1/2 (Network Termination)
  - NT1 terminates the transmission line from the local exchange
  - NT2 handles a variety of functions such as multiplexing and switching
- TE1/2 (Terminal Equipment)
  - TE1 has an ISDN interface and terminates an ISDN call
  - TE2 is similar to TE1 but does not conform to ISDN recommendations
- TA (Terminal Adapter—not required for the Intel Express 8100 Router)

  Allows a TE2 to be connected to the ISDN. It performs a conversion between the signalling and user information formats at the R interface and the ISDN interface. R interfaces are typically V.24 and X.21 interfaces.

Chapter 5

X.25 Services

In this chapter

X.25 is available in the Express 8100 Router with X.25

The WAN port must be configured as an X.25 port before it can be used for X.25 services.

As X.25 uses virtual circuits, a number of links (virtual circuits) can be made via the same X.25 port.

This chapter gives an introduction to X.25 services in the Intel Express 8100 Router.

X.25 Characteristics

Introduction to X.25

X.25 is an internationally agreed protocol for communications between a LAN device (for example a router connected via a modem) and the Packet Switched Data Network (PSDN) node. The packet switched data network can be either a Public Data Network (PDN) or a privately owned network. X.25 is independent of the internal protocols used within the Packet Switched Data Network (PSDN). X.25 uses the OSI standard reference model, and defines standards for layers 1, the physical layer; layer 2, the data link layer and layer 3, the network layer.

Packet Switched Data Networks (PSDNs)

Packet Switched Data Networks work on a store-and-forward principle; data packets received at a node in the PSDN are stored until it is convenient for the network to forward them. This is a cost-effective arrangement as PDNs need not reserve network bandwidth for packet switched connections; packets can be forwarded when the bandwidth becomes available. PSDNs can therefore offer competitive rates for data transfer.

X.25 device characterization

At the physical layer (layer 1 of the OSI reference model), X.25 defines LAN communications devices (for example a router) as Data Terminal Equipment (DTE) and PSDN nodes as Data Communication-terminating Equipment (DCE).

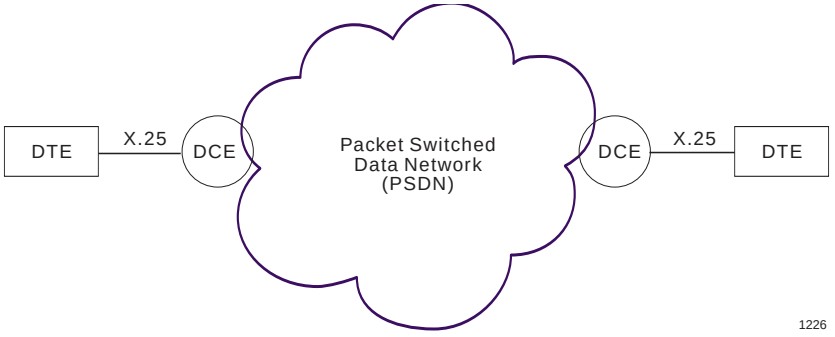

The X.25 protocol defines the communication interactions between the DTE and DCE devices, and is independent of the structure and protocols used between the two DCE devices and over the PSDN.

Physical layer

X.25 specifies the use of full duplex, point-to-point synchronous circuits for the physical layer.

Link layer

The primary responsibility of the link layer protocol is to ensure the error-free exchange of data over a link which has variable delays. The X.25 standard defines the use of either Link Access Procedure (LAP) or Link Access Procedure, Balanced (LAPB) where LAPB is the preferred link layer protocol and is the protocol implemented in the Intel Express 8100 Router.

Network layer

X.25 network layer (also known as packet layer) services give X.25 its virtual circuit characteristics. In addition to sending and receiving data packets, the network layer is responsible for:

- establishing and clearing virtual circuits across the PSDN
- recovering from error conditions (reset and restart procedures)

X.25 Services and Features

Service types available

X.25 allows two basic service access types to a PSDN:

Switched Virtual Circuits (SVCs)

An SVC is a temporary association between two DTEs. This requires that the link between the DTEs through the PSDN must first be established by making what is known as a Virtual Call. SVCs are therefore only initiated by a DTE when data requires transmitting, and are cleared again when the data has been sent. An SVC is analogous to a dial-up (switched) WAN link between two LANs.

Permanent Virtual Circuits (PVCs)

A PVC is a permanent association between two DTEs; no Virtual Call is required to initialize the link and the link is not cleared when the data has been sent. A PVC is analogous to a leased line WAN link.

Logical channels

X.25 uses logical channels when communicating over PSDNs. This means that a number of logical links (connections) to different remote destinations can be established via one physical X.25 port. For example, an X.25 port (one number) with 6 logical channels can communicate with 6 remote destinations simultaneously.

Line speed matching

The store-and-forward nature of PSDNs means that X.25 requires advanced data flow control and line speed matching. This means that data can be received at a different speed than it was transmitted. For example, a host at a central site can be connected via a 56 kbps X.25 connection to several remote sites which are connected with cheaper 19.2 kbps X.25 connections.

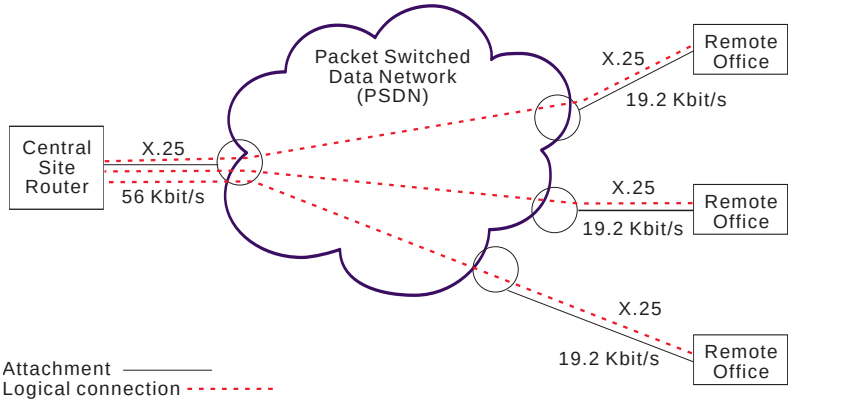

X.25 Packets and Virtual Calls

Supervisory packets

In addition to data packets, X.25 uses special supervisory packets to:

- establish Virtual Calls
- clear Virtual Calls
- negotiate packet and window sizes
- maintain and supervise logical channels across the PSDN
- recover from error conditions

Many of these X.25 supervisory packets are only used by SVCs.

Virtual Call process

Virtual Calls are made using supervisory packets to establish virtual circuits through the PSDN. The Virtual Call process is as follows:

1. The DTE which wishes to establish the Virtual Call sends a Call Request Packet to the local DCE on the highest free Logical Channel Identifier it has available for outgoing calls.
   The Call Request Packet contains the address of the remote DTE.
2. The PSDN decides if the call is valid; if it is, the DCE forwards the packet to the appropriate remote DTE via the remote DCE, on the lowest Logical Channel Identifier available at the remote device.

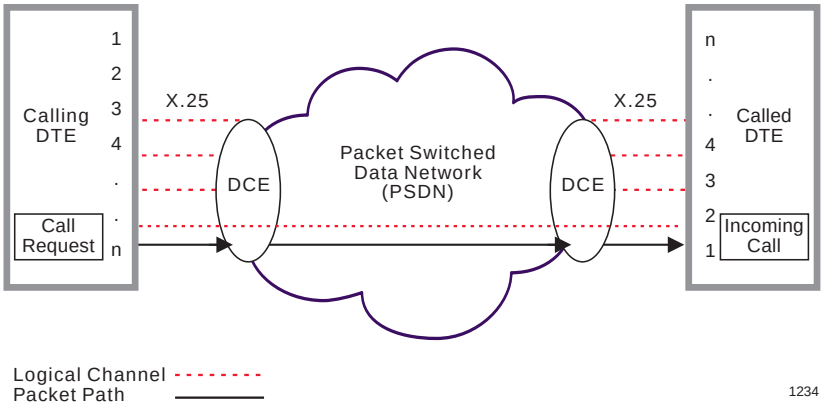

3. If the remote DTE accepts the call, it sends a Call Accepted Packet to the DCE (the remote DCE).
   The Call Accepted Packet contains the Logical Channel Identifier on which the call request was received. The association between the Logical Channel Identifier and the route to the calling DTE has already been established.

4. The remote DCE forwards the call Accepted Packet across the PSDN to the local DCE which then passes it to the calling DTE. Again, only the Logical Channel Identifier from the local DCE to the calling DTE is included in the packet.

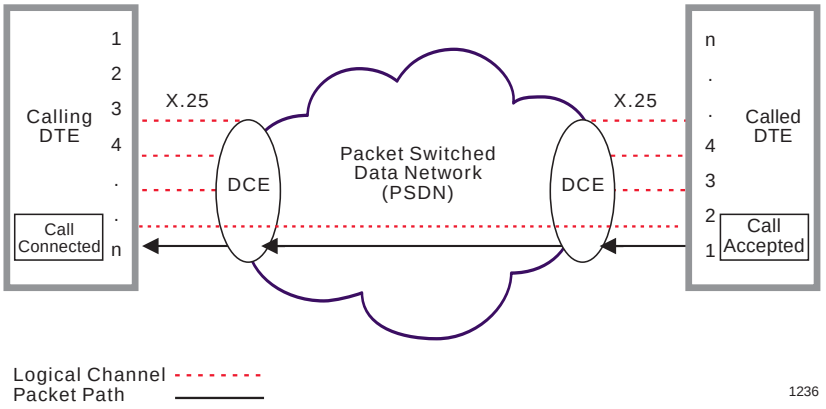

Window size

The window is the number of packets that may be transmitted without receiving acknowledgment.

As X.25 is based upon LAPB which is resilient to packet loss, the window size is used for controlling the flow of packets.

Packet and window sizes

For PVCs, the packet and window sizes used when communicating across the X.25 network are fixed (static) and are defined by the X.25 service provider (typically packet size of 128 bytes and a window size of 2 packets).

For SVCs, some X.25 networks allow negotiation of the packet and window sizes between the DTE and the DCE. The packet and window sizes used are arrived upon using an algorithm defined by the X.25 standard. Parameter negotiations can give more efficient data transfers but be aware that some X.25 network providers charge extra for the negotiation services.

For SVCs across X.25 networks or subscriptions which do not allow negotiation, the packet and window sizes are fixed (static) and are defined by the X.25 service provider (typically 128 bytes, 2 packets).

Note The packet and window sizes are only used between a DTE and the DCE. The packet and window size used at the remote DTE and DCE can be different.

Encapsulation

Data packets for transmission over an X.25 link must be encapsulated. The Intel Express 8100 Router supports the following encapsulation types:

Single-protocol Encapsulation of Datagrams

The Intel Express 8100 Router allows single-protocol encapsulation of IP or IPX datagrams for transmission over an X.25 link as specified in RFC 1356. Single-protocol encapsulation of IP datagrams also conforms to RFC 877 which is an older standard which only allows the encapsulation of IP datagrams.
When only IP or IPX packets are to be transmitted over an X.25 link, the single-protocol encapsulation modes reduce the transmission overheads.

Multiprotocol Encapsulation of Datagrams

The Intel Express 8100 Router allows multiprotocol encapsulation of datagrams for transmission over an X.25 link as specified in RFC 1356. If both IP and IPX are to be transmitted, multiprotocol transmission must be used. Multiprotocol encapsulation must also be used if bridging is to be supported over the X.25 link.
For multiprotocol encapsulation:

- IP packets are encapsulated according to RFC 1356 and are identified using a Hex CC in the header
  RFC 1356 also allows SNAP encapsulation of IP packets which is not supported by the Intel Express 8100 Router.
- IPX packets are SNAP 8137 encapsulated according to RFC 1356

X.25 Addressing

X.121

X.25 uses CCITT recommendation X.121 “International Numbering Plan for Public Data Networks” for addressing within international public data networks. This recommendation specifies how data network addressing (numbering) can be implemented along side existing international telephone networks.

X.121 number format

An entire X.121 number (address) is known as an International Data Number and consists of from 5 up to 14 digits. The International Data Number can comprise the following:

Data Network Identification Code (DNIC)

This consists of 4 digits for international data network addressing. The first 3 digits are the Data Country Code (DCC) and identify a specific country. The fourth digit is known as the network digit and is used to distinguish between different data network providers (public or private) operating within the country.
Assignment of DCCs is the responsibility of the ITU (previously called the CCITT). The network digit is assigned nationally.

Network Terminal Number (NTN)

Network addresses within a DNIC (a particular data network provider within a particular country) are called Network Terminal Numbers (NTNs). NTNs identify a DTE/DCE interface and are assigned by the data network provider. NTNs can be of variable length up to 10 digits, depending on the network provider.

National Number (NN)

Where, for example, an integrated numbering scheme exists within a country and different network providers are identified within the national numbering scheme, National Numbers (NNs) can be used instead of Network Terminal Numbers (NTNs). As National Numbers already identify a data network provider within the country, the International Data Number therefore consists of the Data County Code (DCC) plus the National Number (NN).
Therefore, a International Data Number = DNIC + NTN or DCC + NN.

Sub-addressing

This can consist of any remaining digits of the 14 digit X.121 number, and in some countries it can be used for sub-addressing within the International Data Number, for addressing a number of devices/nodes located at a single International Data Number. This can be used for example, for addressing different links in the router.

Address checking

To ensure a one-to-one relationship between end devices, a unique number is assigned to each link across the PSDN. Incoming calls are only accepted if the local and remote addresses match, as shown in the table below:

Addresses Must Match for a Call to be Accepted	
Calling Device	**Receiving Device**
Local Address	Remote Address
Remote Address	Local Address

Destination address check

The router only accepts an incoming call when the `Remote Address` matches the `Local Address` of an X.25 link in the router.

If a `Local Address` is not defined for an X.25 link, all incoming addresses are accepted by the link (the address check is not performed).

Source address check

Source address checking adds security to the WAN link—the router only accepts an incoming call when the source address matches the `Remote Address` configured for the X.25 link.

Chapter 6

PPP Multilinks

In this chapter

This chapter gives an introduction to PPP Multilinks in the Intel Express 8100 Router.

Introduction

The ISDN Intel Express 8100 Routers implements PPP Multilink facilities which allow multiple PPP links between two sites to be used as a single route. A PPP Multilink allows links to be added dynamically (for bandwidth-on-demand) or statically and has almost the combined bandwidth of the individual links.

Channel bundling

Combining multiple links into a single route (link) is also sometimes known as channel bundling.

Master link

The master link is the PPP Multilink between two sites and consists of the slave links that make up the Multilink. The master link is the controlling link between the two sites, and controls (for example) when the multilink is established or not, data encryption over the link, peer authentication over the link using CHAP, and so on.

Slave links

Slave links are individual PPP links that are defined as part of the master link. These links are ISDN links which correspond to ISDN B-channel links.

Configuration overview

The following shows an example of how Multilink is set up in the router:

<table>
<tr><td colspan="3">ISDN port</td></tr>
<tr><td colspan="2">ISDN slave link</td><td>ISDN slave link</td></tr>
<tr><td colspan="3">PPP Multilink</td></tr>
<tr><td>IP Routing</td><td>IPX Routing</td><td>WAN Bridging</td></tr>
</table>

The ISDN port on the router must first be configured, the Multilink is configured to combine ISDN slave links, and then the routing protocols (IP and Novell IPX) and bridging are configured on top of the Multilink as required.

Operating modes

The slave links that make a PPP Multilink can be used in two basic operating modes:

Bandwidth-on-Demand (BOD) Mode

When a slave link of a PPP Multilink is used in Bandwidth-on-Demand mode, the link is dynamically allocated when the data traffic volume requires its use, and closed down when not required. This facility offers the most cost-effective use of dial-up PPP links (including ISDN B-channels) where the cost of operation depend on the time that the link is in use.

Bandwidth-on-Demand can only be done when a slave link is a dial-up (switched) link such as ISDN B-channels. Leased line PPP links can only be statically allocated, that is they are always available for use by the PPP Multilink.

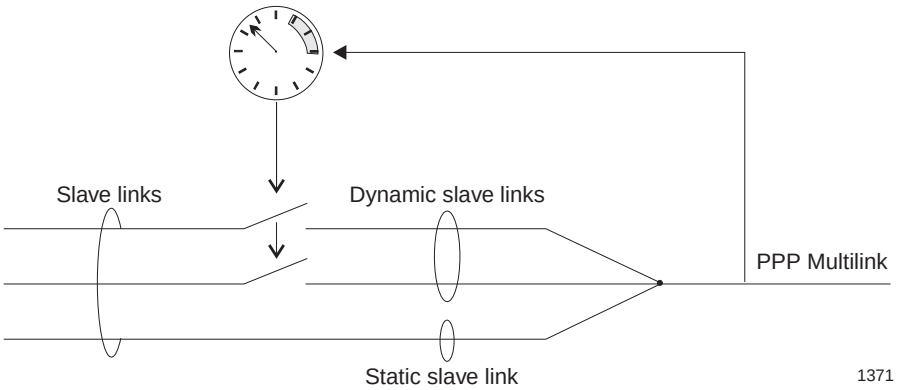

The criteria used to open and close Bandwidth-on-Demand links can be:

- on the amount data awaiting transmission
- on the amount of received data
- both the amount of data awaiting transmission and the amount of received data

Static mode

When a slave link of a PPP Multilink is used in Static mode, the link is always available for use by the PPP Multilink.

Internet Tunnels

Internet Tunnels

Internet Tunnels are not a direct link via a router port but are secure links established via the best IP route to an Internet Service Provider (ISP).

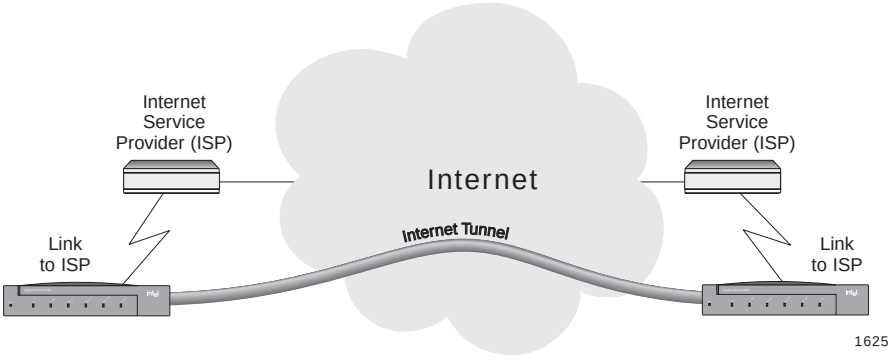

Internet Tunnels are a cost effective solution to transferring data between remote sites; IP and IPX routing and bridging are available over the tunnel as with a conventional WAN links.

The costs of operation of an Internet Tunnel can be very low when compared to routing via other WAN link types an ISDN link, especially when transmitting data over long distances. The throughput of the Internet Tunnel is dependent on the Internet so can be quite slow at certain times of the day.

Note It is strongly recommended that you use the Setup Wizard in Intel Device View for Windows to set up your first Internet Tunnel. If other Internet Tunnels require configuring, use the tunnel configured by the Setup Wizard as a guide.

Internet Tunnel considerations

The following considerations must be taken into account when setting up an Internet Tunnel:

IP Routing using RIP over the Internet Tunnel

When using IP RIP routing over an Internet Tunnel a static host route to the remote router via the link from the router to the Internet Service Provider must be established.

If RIP routing is used without the static host route, the router tries to use the Internet Tunnel as the best route to the remote network. The static host route ensures that the Internet Tunnel is established via the route to the ISP and that all the data to the remote router is transferred through the Internet Tunnel.

Site Access Security

Access security must be considered when using Internet Tunnels. This can be achieved using IP filtering to eliminate unauthorized traffic from the link.

Hiding your internal LAN address from the Internet

Dynamic Network Address Translation (NAT) offers the best solution for hiding your internal LAN from the Internet. See “Network Address Translation (NAT)”, p. 57.

Data Security using Encryption

The use of data encryption over Internet Tunnels is highly recommended. Private data being transferred over the public Internet should always be encrypted for security.

Note Data encryption is only available in certain models of the Intel Express 8100 Router which are available in some countries.

The protocol structure of an Internet Tunnel

Internet Tunnels have the following protocol stricture:

IP routing	IPX routing	Bridging
PPP		
TCP		
IP static host route		

IP routing services then use the best link to the remote IP network (the link via the Internet Service Provider). For example, this could be an ISDN link, in which case the full structure would be

IP routing	IPX routing	Bridging
PPP		
TCP		
IP static host route		
PPP		
ISDN port		

Data security over X.25 links application of Internet Tunnels

To allow data security over X.25 links, Internet Tunnels can be used on top of them. this offers all the security features supported by PPP (data encryption, CHAP and PAP) on the X.25 links.

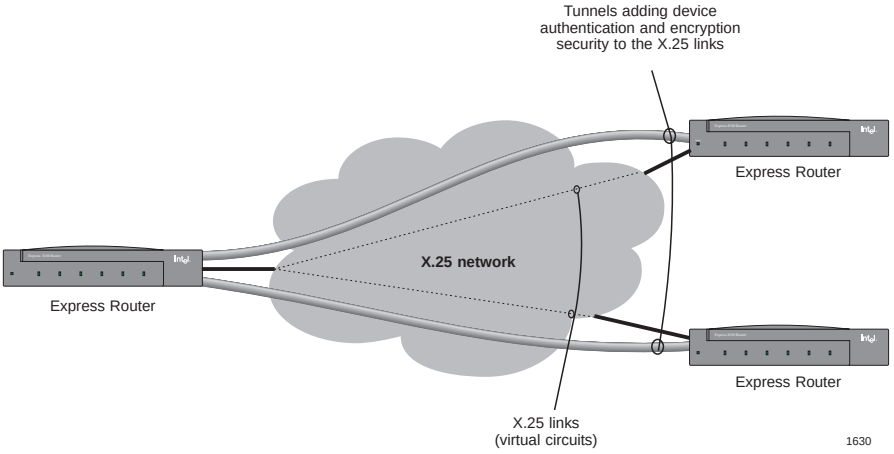

To establish a tunnel over an X.25 link, the X.25 link is set up as normal, together with an Internet Tunnel to the IP address of the remote router over the X.25 link. Data security is configured on the tunnel as required. A single static route is then

configured on the X.25 link to the IP address of the remote router (see “Static Routes”, p. 68). Routing and bridging for data over the link are then configured on the Internet Tunnel.

<table>
<tr><td colspan="3">X.25 port (WAN port configured for X.25)</td><td rowspan="5">- Up to 2 links -</td></tr>
<tr><td colspan="3">X.25 link (PVC or SVC)</td></tr>
<tr><td colspan="3">Static IP route to the IP address of the router over the X.25 link</td></tr>
<tr><td colspan="3">Internet Tunnel</td></tr>
<tr><td>IP Routing</td><td>IPX Routing</td><td>WAN Bridging</td></tr>
</table>

Part II

IP, Novell IPX and Bridging Services

Chapter 8

IP Routing

In this chapter

This chapter describes IP Routing services based upon the Routing Information Protocol (RIP—both RIP-1 and RIP-2 are supported), and using static routes implemented in the Intel Express 8100 Router.

IP Concept

Concept description

In an IP environment, any station connected to an IP network is referred to as a host. IP Routing is based upon RIP (Routing Information Protocol). When IP packets are routed, they are routed from a source host to a destination network without knowing if the destination address exists on the destination network.

Background knowledge

It is assumed that you are familiar with the terminology and the protocols used in IP environments.

References

See the Preface for a list of Requests for Comments (RFCs) concerning IP routing.

Address assignment

An IP address consists of a network part and a host part. The network part of the IP address must be globally unique and are assigned by InterNIC (International Network Information Center). However, addresses are often provided by the network provider. For more information about the InterNIC project consult the RFC 1594.

The host part of an IP address is the responsibility of the network manager.

Private networks

In private networks, where connections to other IP networks are not desired, locally assigned network addresses can be used.

BootP broadcasts only

In the current IP Routing implementation in the Intel Express 8100 Routers the only broadcast packets forwarded are BootP broadcasts.

Frame types and type codes

IP Frame types are described in Appendix D “IP Frame Formats”, p. 126. The following Ethernet type codes are used in IP environment, also see Appendix C “Ethernet Type Codes”, p. 121:

Type field	Description
0800	DOD Internet Protocol (IP)
0806	Address Resolution Protocol

IP Addresses Structure

Address notation

IP addresses are 32-bit numbers. The most common notation for IP addresses divides the 32-bit address into four 8-bit fields and specifies the value of each field as a decimal number (from `0` to `255`, each representing an 8-bit octet). Each number field is separated by a period (for example 14.0.65.3). This is called the dotted decimal notation.

Network numbers

The 32-bit address field consists of a network and a local host part. They are divided into different address classes which differ in the number of bits allocated to the network part and the hosts part (local address) of the address. The value of the first octet in the IP address defines the address class (classes A, B, C, D).

Class A address

The class A address comprises a 7-bit network number and a 24-bit local address. The highest order bit is set to `0`. This allows 126 class A networks.

1								2								3								4							
7	6	5	4	3	2	1	0	7	6	5	4	3	2	1	0	7	6	5	4	3	2	1	0	7	6	5	4	3	2	1	0
0	Network							Local Address																							

Class B address

The class B address comprises a 14-bit network number and a 16-bit local address. The two highest-order bits are set to `1 0`. This allows 16,256 class B networks.

1								2								3								4							
7	6	5	4	3	2	1	0	7	6	5	4	3	2	1	0	7	6	5	4	3	2	1	0	7	6	5	4	3	2	1	0
10	Network															Local Address															

Class C address

The class C address comprises a 21-bit network number and a 8-bit local address. The three highest-order bits are set to `1 1 0`. This allows 2,072,640 class C networks.

1								2								3								4							
7	6	5	4	3	2	1	0	7	6	5	4	3	2	1	0	7	6	5	4	3	2	1	0	7	6	5	4	3	2	1	0
1 1 0								Network																Local Address							

Class D address

The class D address is used as a multicast address. The four highest-order bits are set to `1 1 1 0`.

1								2								3								4							
7	6	5	4	3	2	1	0	7	6	5	4	3	2	1	0	7	6	5	4	3	2	1	0	7	6	5	4	3	2	1	0
1 1 1 0				Multicast Address																											

Note No addresses are allowed which have the four highest-order bits set to 1 1 1 1 (also known as class E address).

Addresses available

The following IP addresses are available for the different IP address classes:

Note n = network part of the address,
h = host part of the address.

Class	Address Range available	Notation
A	1.0.0.0 through 126.0.0.0	n.h.h.h
B	128.0.0.0 through 191.254.0.0	n.n.h.h
C	192.0.0.0 through 223.255.254.0	n.n.n.h
D	224.0.0.0 through 239.255.255.255 for multicasts.	n.n.n.n
E	240.0.0.0 through 247.255.255.255 reserved.	n.n.n.n

IP Subnets

Why subnetting?

The one-level network addressing scheme and the limited number of IP addresses impose limitations on network complexity. Therefore subnetting has been introduced.

What is subnetting?

Subnetting is the process of dividing an IP network with a single IP address into two or more smaller IP networks (subnetworks). This involves dividing the host part of the IP address into a subnet number and a host number.

Characteristics

An IP subnet functions as an independent network. To a remote network it appears to be a single discrete network.

How to create subnets

To create a subnet you must impose a hierarchy on host addresses which allows for multiple subnets to exist within a single class A, B or C network. It is the high order bits of the hosts part that specify the subnet. Derive the IP network by ANDing the IP address with the subnet mask:

1								2								3								4							
7	6	5	4	3	2	1	0	7	6	5	4	3	2	1	0	7	6	5	4	3	2	1	0	7	6	5	4	3	2	1	0
Class B address (subnet mask 255.255.0.0)																															
Network part																Local host part															
Class B subnet (subnet mask for example 255.255.192.0)																															
Network part																		Local host part													

An example of the network part looks like this:
Decimal Notation: `16.1.0.1` AND `255.255.192.0` = `16.1.0.0`
Hexadecimal Notation: `0x10010001` AND `0xffffc000` = `0x10010000`

Subnet masks

A subnet mask specifies which part of the host's address is used to define the subnet. It is a 32-bit IP address written in dotted-decimal notation with all ones (binary) in the network and subnet part of the IP address.

Note When a subnet mask is used on an IP address, the resulting host part of the address must not end in zero or 255.

Structure of a subnet mask

The following table shows the structure of a subnet mask, where the field containing a 1 is the subnet part and the field containing a 0 is the local hosts part of the address:

Address (bit no.)									Subnet Mask	
7	6	5	4	3	2	1	0		hexadecimal	decimal
0	0	0	0	0	0	0	0		0	0
1	0	0	0	0	0	0	0		80	128
1	1	0	0	0	0	0	0		C0	192
1	1	1	0	0	0	0	0		E0	224
1	1	1	1	0	0	0	0		F0	240
1	1	1	1	1	0	0	0		F8	248
1	1	1	1	1	1	0	0		FC	252
1	1	1	1	1	1	1	0		FE	254
1	1	1	1	1	1	1	1		FF	255

Example

An example looks like this:

Class C address	193.88.251.1
Class C network	193.88.251.
Class C host	.1
Divide the network by two more bits, that is set the subnet mask to 255.255.255.192 to obtain 4 subnetworks:	
Subnetworks	193.88.251.0 available for hosts 1 to 62 193.88.251.64 available for hosts: 65 to 126 193.88.251.128 available for hosts: 129 to 190 193.88.251.192 available for hosts: 193 to 254

Address Resolution

Address resolution

The process of determining the MAC address (48-bit) from an IP address (32-bit) is called Address Resolution (AR).

Address resolution with ARP

The Address Resolution Protocol (ARP) is used for address resolution of hosts which have not previously been communicated with (for example, hosts which have just been switched on and added to the network). When the Intel Express 8100 Router receives an IP data packet which has to be transmitted to a local un-

known host, it broadcasts an ARP request packet to all local hosts. Only an active hosts with the specified IP address reply with an ARP reply packet that contains its MAC address. The packet exchange with this host is then initiated.

Address resolution with Proxy ARP

Proxy ARP is used for address resolution of remote hosts situated on another network that a local host wants to communicate with. When an ARP request is received from a remote host, the router having the best route to it sends an ARP reply packet with its own MAC address to the originator. The ARP request originator then sends its packet to that particular router to initiate the packet transmission to the destination host. Proxy ARP is used if the network contains any hosts that do not allow the configuration of a primary router.

Internet Control Messages

Introduction

To control and manage IP connections, the following tools are available:

- ICMP (Internet Control Message Protocol)
- PING function (IP Packet InterNet Groper function)

ICMP protocol

The ICMP (Internet Control Message Protocol) is a protocol that focuses on control and management of IP connections. ICMP messages are generated by routers detecting a problem within the IP part of the packet header. A router may alert other routers or inform source/destination hosts using ICMP messages.

Ping function

The IP Packet InterNet Groper Function (PING) initiates the transmission of ICMP Echo messages by the router. The PING function can be initiated from Intel Device View for Windows.

How ICMP and PING work

A router that receives an ICMP Echo message which was initiated by the PING function sends an ICMP Echo Reply message to the ICMP Echo message source.

IP Filters—Firewall Protection

Filtering in the Intel Express 8100 Router

The Intel Express 8100 Router implements comprehensive IP filters to restrict access between networks (firewall protection) and to reduce unnecessary internetwork data traffic to save bandwidth.

Note A good knowledge of IP, the IP protocols (TCP, ICMP and UDP) and the use of TCP and UDP ports is required before IP filtering can be set up effectively.

IP packets filtered

IP traffic to and from the router can consist of packets being forwarded (routed) via the LAN and WANs and IP packets to the router itself (RIP, TELNET, SNMP, TFTP, and so on). Both of these kinds of IP packets can be discarded by the filters.

Note In order for the router to operate correctly, the necessary packets must be allowed to pass by the filters.

Transmit and receive filters

Filters are defined on a link basis, and separate filters are implemented for transmitting (for restricting IP packets leaving the router on a particular link), and receiving (for restricting IP packets entering the router from a particular link).

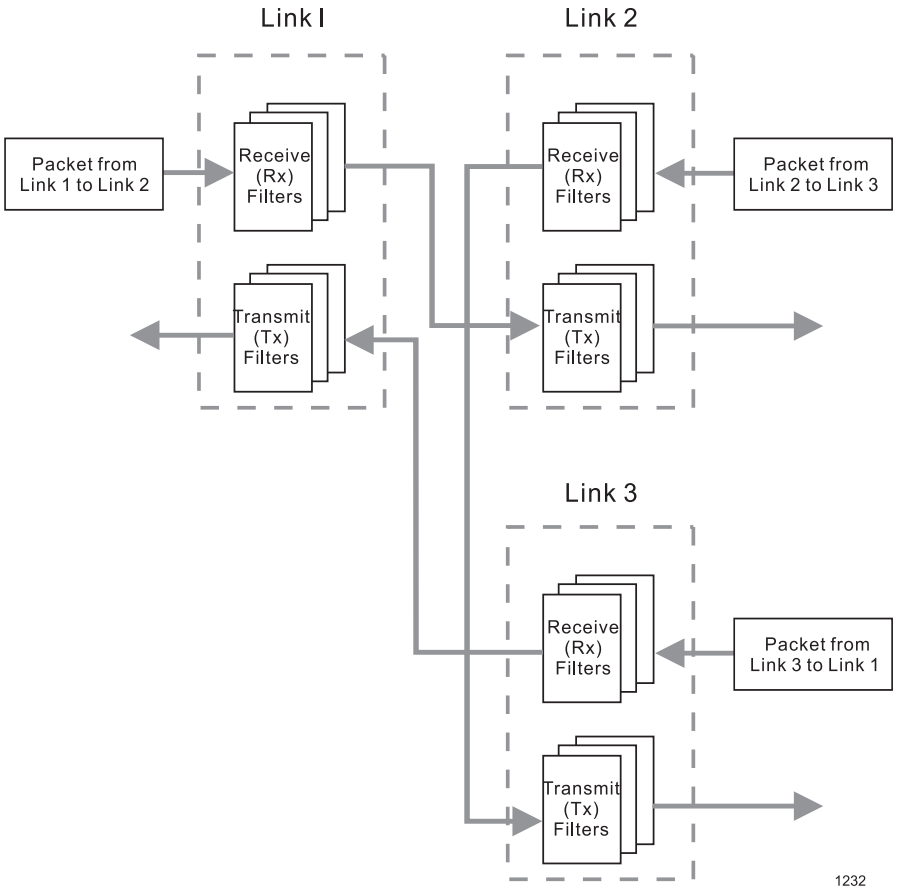

The above diagram shows how packets are processed by the Receive and Transmit filters implemented for different links.

Note Filters must be defined as strictly as possible to eliminate unauthorized access to services, hosts and networks.

Filter criteria

Each filter in the Intel Express 8100 Router can be setup to pass or discard IP packets based on the following criteria:

IP Protocol

A filter can process packets based on all IP protocols, User Data Protocol (UDP), Transmission Control Protocol (TCP). Filtering can also be based on all TCP flags or just with the acknowledge (ACK) flag set, Internet Control Message Protocol (ICMP) or another protocol defined by the protocol number.

Source Address

To filter packets entering the router via the link from a specific host or network.

Source Port

To filter packets originating from a single port (for example FTP data, TELNET, SMTP), from a range of ports or all ports.

Destination Address

A filter can process packets addressed to a host address or a network address.

Destination Port

A filter can process packets addressed to a single port (for example FTP control, TFTP, SNMP), a range of ports or all ports.

Up to 50 IP filters can be defined for all the IP links.

References for Firewall Protection

Defining filters for firewall protection to ensure unauthorized access to services, networks and hosts can be very complex. Two very good books on the subject are:

- Building Internet Firewalls
  D. Brent Chapman and Elizabeth D. Zwickey
- Firewalls and Internet Security

  Bill Cheswick and Steve Bellowin

Filtering process

The Intel Express 8100 Router can contain a list of filters for each link for both transmitting and receiving data via the link. These filters are used to filter IP packets as follows:

1. When enabled, outgoing IP packets on the link are checked against the Tx filters sequentially (using the first filter defined, then the second, and so on). Similarly, incoming IP packets from a link are checked against the Rx Filters.
2. If a filter is found where all the filter criteria match those of the packet, it is used to pass or discard the packet as setup for the filter. The rest of the filters are then ignored.

3. If no filter entry is found where all the filter criteria match those of the packet, the default filter action is used to either pass or discard the packet.

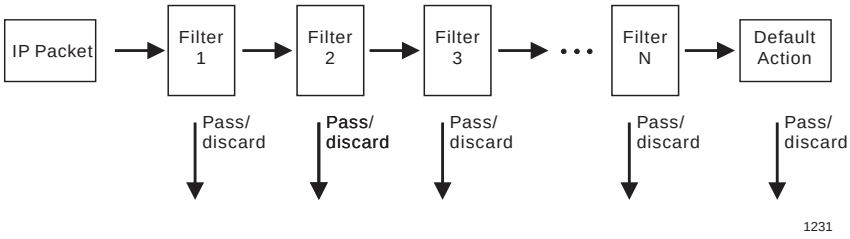

Logging filtered packets

The router allows filtered packets to be logged. This facility can be configured for both packets filtered by filter entries and packets filtered by the default filter action. This is a useful tool for troubleshooting the filters to ensure that they are operating as expected, and for detecting hits on the filters in case of an attack from an unauthorized source.

Note Logging reduces the performance of the router slightly and should therefore only be enabled when required—when troubleshooting or when an attack is suspected.

For packets which are logged, the packet details are recorded in the System Log for the router which cabn be displayed for Intel Device View for Windows. The details logged are the source and destination address and port and the IP protocol.

IP Filter Example

Filter requirement

Host 1 on LAN 1 must have TELNET access to remote Host 2 on LAN 2 (but Host 2 must not be able to establish a TELNET session onto Host 1).

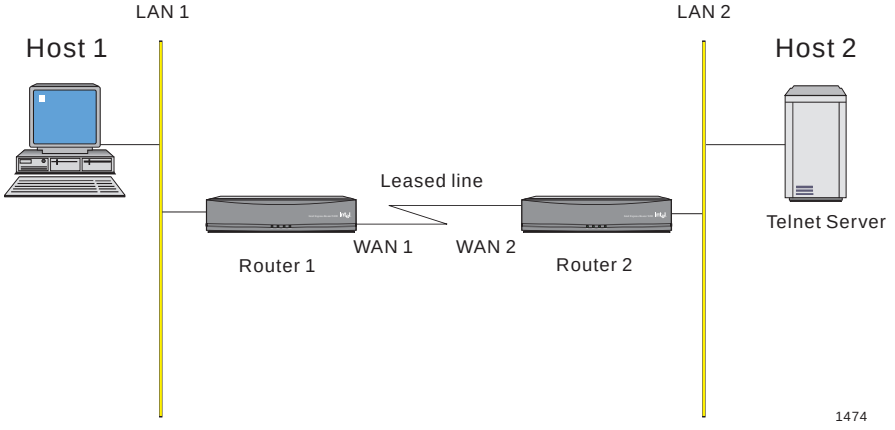

IP traffic

The TCP/IP traffic to establish a TELNET session from Host 1 to Host 2 is:

Direction	Flags	Source Add.	Source Port	Dest.	Dest. Port	Comments
1 > 2	ACK = 0	Host 1	X	Host 2	23	TELNET request from Host 1 to Host 2.
2 > 1	ACK = 1	Host 2	23	Host 1	X	TELNET response from Host 1 to Host 2—connect to Port X (where X > 1023).
1 > 2	ACK = 1	Host 1	X	Host 2	23	TELNET command.

Filters Required

Filters are required in Intel Express 8100 Router 1 to protect LAN 1 from unauthorized access and so that only Host 1 is allowed to establish a TELNET session onto Host 2.

The filters required are as follows:

Rx Filter on WAN 1 of Intel Express 8100 Router 1

This filter only allows TCP access from the WAN link from Host 2 to Host 1, with the ACK flag set (session established), a TELNET Server source port and a destination port of greater than 1023 (unassigned ports used by TELNET Client).

Parameter	Setting
Default Filter Action (on Advanced screen)	Discard
Action (Filter action)	Pass
Protocol	TCP
TCP Flags	ACK
Src. Address Type	Host
Src. Address	Host 2’s IP address
Src. Port	TELNET
Src. Port Operator	==
Dest. Address Type	Host
Dest. Address	Host 1’s IP address
Dest. Port	Other
Dest. Port Value	1023
Dest. Port Operator	>

Tx Filter on WAN 1 of Intel Express 8100 Router 1

This filter only allows TCP access to the WAN link from Host 1 to Host 2, with a source port of greater than 1023 (unassigned ports used by TELNET Client) and a TELNET Server destination port.

Parameter	Setting
Default Filter Action (on Advanced screen)	Discard
Action (Filter action)	Pass
Protocol	TCP
TCP Flags	All
Src. Address Type	Host
Src. Address	Host 1’s IP address
Src. Port	All
Src Port Value	1023
Src. Port Operator	>
Dest. Address Type	Host
Dest. Address	Host 2’s IP address
Dest. Port	TELNET
Dest. Port Operator	==

Note RIP updates are also filtered out by these filters, so dynamic routing (using RIP) cannot function. This can be overcome by defining additional filters to allow RIP updates to pass. Alternatively, a static route to LAN 2 can be established in Router 1 and a static route to LAN 1 can be established in Router 2.

Network Address Translation (NAT)

Introduction to NAT

Network Address Translation (NAT) can be used to translate local (internal) IP network addresses to remote (external) IP addresses used by devices outside the internal network. For example, this can be done to connect a site which uses an unreserved network address to a site which requires a reserved network address assigned by InterNIC (see “Address assignment”, p. 47), such as the Internet. NAT can also be used for connecting multiple sites which use the same IP network address.

Mapping methods

NAT can use one of the following mapping methods:

Static Mapping

Static Mapping can be used to translate an internal network to an equal sized external network (for example a class B network to a class B network). Static Mapping, can also be used to translate individual internal IP addresses to external IP addresses.
Static Mapping is described in "Static Mapping", p. 60.

Dynamic Mapping

Dynamic Mapping can be used to translate between IP networks of different sizes, that is, a large internal network can be translated to smaller external network or vice versa (for example a class B internal network could be translated to class C external network addresses).
Dynamic Mapping is described in "Dynamic Mapping", p. 60.

How NAT works

NAT uses a Network Address Translation table which contains internal and corresponding external IP addresses. Each entry in the table can be either a network address, a subnet address or an individual IP address.
For packets from the router, addresses matching entries in the internal list are translated to the associated external addresses. For packets from the remote router, addresses matching entries in the external list are translated to the associated internal addresses.

Note The internal and external addresses are NOT source and destination IP addresses, but are addresses which are translated if they pass through the router.

When an IP packet is translated by NAT, both the source and destination IP addresses can be translated by different entries in the table, as shown in the following diagram.

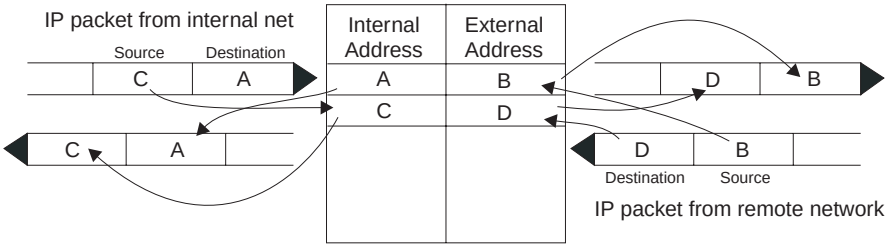

1351

For packets from the internal network, address A for either the source or destination IP address is translated to address B, and address C is translated to address D. For packets from the external network, address B for either the source or destination address is translated to address A, and IP address D is translated to address C.

NAT and routing

For IP packets transmitted from the router, the packet is first routed to the appropriate LAN or WAN link using the destination address included in the packet from the source then translated using the NAT table for that link.

For packets from the external network, the IP addresses are translated as they enter the router. If they are then transmitted on another link on which NAT is enabled, the packet is first routed to the appropriate LAN or WAN link using the destination address (which may already have been translated) then translated again using the NAT table for the outgoing link.

NAT on the LAN link

It is possible to define NAT translation tables for the LAN link but it may be confusing to think of the LAN as an external network. NAT on the LAN link works in exactly the same way as for the WAN links.

NAT and IP filtering

When NAT is used in conjunction with IP filtering, for IP packets received by the router, the packet addresses are first translated by NAT then the packets are filtered by the IP Rx filters. For IP packets transmitted from the router, the packets are first filtered by the IP Tx filters then the packet addresses are translated by NAT.

NAT and RIP updates

When NAT is used on a link, RIP updates received or sent via the link are discarded if they do not match a translation entry in the NAT table. To pass all unmatched RIP updates instead of discarding them, a static entry can be defined with internal and external IP addresses of `0.0.0.0` and a network mask of `0.0.0.0`.

Static Mapping

Static Mapping of network addresses

When using Static Mapping of network addresses, addresses are simply converted by translating the network part of the IP address between the internal and external address. The host part of the address remains the same, for example, an internal class B network address 10.10.4.8 (where 10.10 is the network part of the address) could be translated to the external class B network address 177.4.4.8.

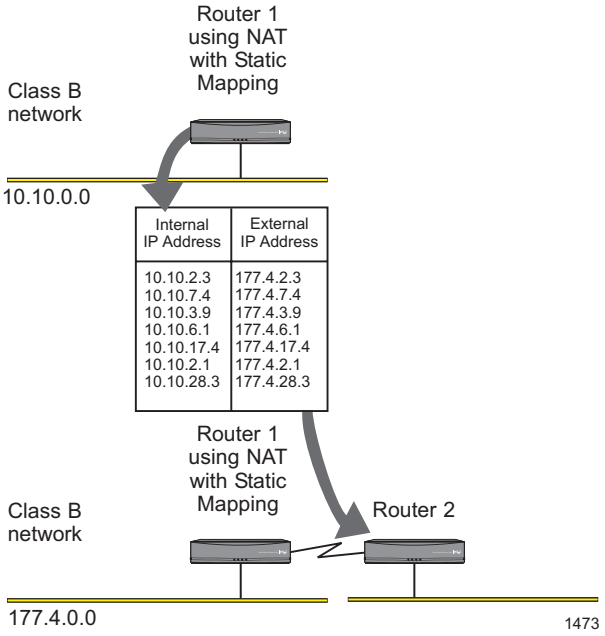

The above diagram gives an example of Static Mapping between two class B network addresses. The internal network address is 10.2.0.0 and the external network address is 177.4.0.0. Only the network part of the address is translated.

Static NAT for individual IP addresses

Entries can be made in the NAT table for individual IP addresses. Internal addresses are simply converted to the corresponding external addresses.
This can be used when Dynamic Mapping is being used, to make internal network devices accessible to the external network.

Dynamic Mapping

Dynamic Mapping

When using Dynamic Mapping, more addresses can be available to either the internal network or the external network. For example, if the internal network is class B and the external network address is class C, the internal network can have up to 65,536 network addresses while the external network address only offers

up to 256 addresses. In this case, the entire internal address (network and host part) must be translated to an assigned external address. External addresses are therefore assigned sequentially as they are required.

Note When using Dynamic Mapping, only internal networks can initiate communications with external devices. External devices do not know the address of the internal device until it has been informed of it by a packet from the device. For this reason, static links must be established to internal network devices which you wish to make accessible to the external network.

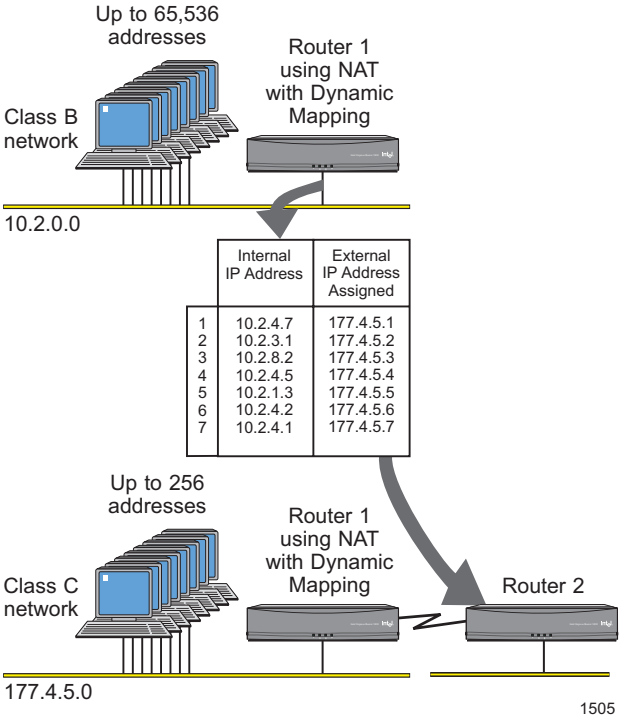

The above diagram gives an example using Dynamic Mapping between an internal class B network (address 10.2.0.0) and an external class C network address (address 177.4.5.0). The full IP address is translated and the internal host addresses are assigned external IP addresses sequentially when they initiate communications over the IP link.

Communications between internal and external addresses using Dynamic Mapping

When using Dynamic Mapping, only internal devices can initiate communications with external devices. Internal addresses are unknown to the external network and a translation address is assigned dynamically whenever an internal device initiates communications with an external device. Following the initial communication, external addresses can then communicate with the translated address for as long as it remains in the NAT table. Internal addresses remain in the NAT table for at least 1 hour following the last communication—that is, the timer

is restarted following every communication. After 1 hour, the address can be reused for another address translation if required but remains available until it is reused for another translation.

Static Mapping to internal addresses for Dynamic Mapping

When using Dynamic Mapping, only internal devices can initiate communications with external devices. To make specific internal addresses permanently accessible to the external network, Static Mapping to these addresses can be established. Static Mapping assign a fixed translation address to an internal device address so that communications can also be initiated to the device from the external network.

Examples of Use

Simple address translation

A common application of NAT is simple address translation between internal and external network addresses. This may be done, for example, to avoid renumbering an entire network which is using an unreserved network number, to be connected to network which requires reserved network numbers (for example the internet) assigned by InterNIC (see “Address assignment”, p. 47).

Address translations between networks of the same size is described in “Static Mapping”, p. 60. Address translations between networks of different size is described in “Dynamic Mapping”, p. 60.

NAT between networks with the same network number

When using NAT, it is possible to have a number of networks which use the same network number. In this case, an artificial network number must be created to which packets are addressed. Two entries must be made in the NAT table to convert both the source and destination IP address fields of packets transmitted between the two networks (see the following example).

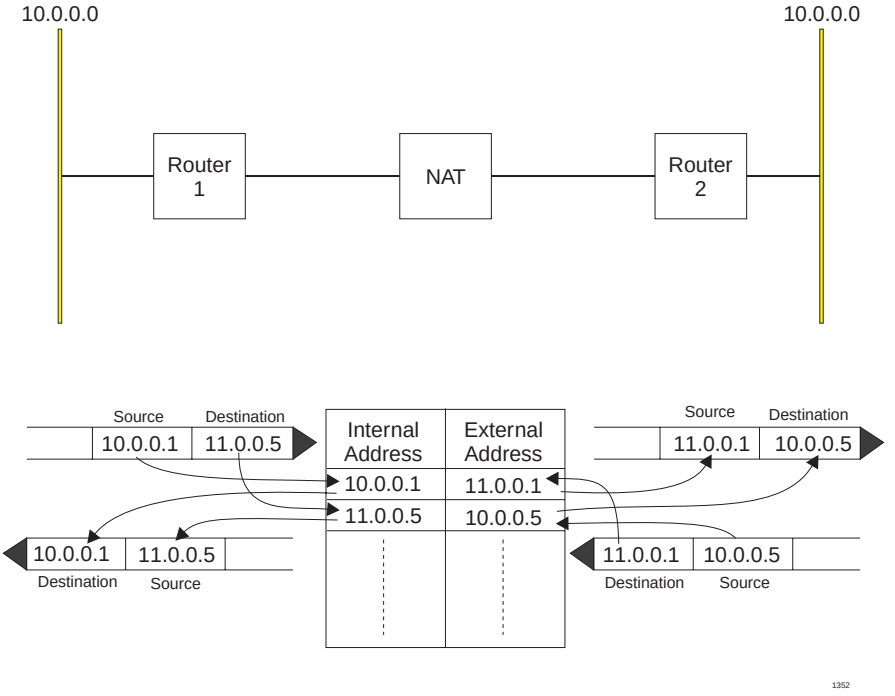

Translation table entries

For the above example, the entries required in the NAT table are:

- Entry 1
  - Mapping: Static
  - Internal Address: 10.0.0.0
  - Internal Mask: 255.0.0.0
  - External Address: 11.0.0.0
- Entry 2
  - Mapping: Static
  - Internal Address: 11.0.0.0
  - Internal Mask: 255.0.0.0
  - External Address: 10.0.0.0

Routing Information Protocol (RIP)

Definition: RIP

Routing information is exchanged by the Routing Information Protocol (RIP), which is an Interior Gateway Protocol (IGP) based on a “Distance Vector Algorithm”. RIP uses the User Datagram Protocol (UDP) to exchange routing information.

RIP-1 and RIP-2

RIP based routing

Routers inform each other about present available paths (router links to other networks) by RIP updates. These are sent periodically (each 30 seconds) or triggered by topology changes. Each router finds the best path to any known remote network. If a router does not receive an update from an adjacent router for 180 seconds or more, it marks the routes served by the adjacent router as unusable.

RIP-1 and RIP-2 standards

The Intel Express 8100 Router supports both RIP-1 and RIP-2. RIP-1 is the original standard as defined by RFC 1058. RIP-2 is an update of RIP-1 and is defined by RFC 1723. RIP-1 and RIP-2 must be considered as separate protocols.

Important differences between RIP-1 and RIP-2

The most important differences between RIP-1 and RIP-2 are summarized as follows:

RIP-1	RIP-2
Network addresses must belong to IP address class A, B or C. Supernetting is not allowed and subnetting is only allowed if the same subnet mask is used throughout the network.	Network addresses are classless—variable length network masks can be used allowing supernetting and subnetting.
Routing updates do not contain subnet masks.	Routing updates contain subnet masks.

RIP-1	RIP-2
Routing updates are sent as broadcasts—most network devices listen to them.	Routing updates are sent as multicasts and many network devices do not listen to them.
Subnet hiding is only done when sending routing updates on numbered links—subnetworks are summarized by routers.	No subnet hiding is done. This results in larger routing tables as all subnetworks are listed.
Routing updates are not authenticated.	Routing updates are authenticated—password protection can be set up to ensure that a routing update cannot be passed to a router from an unauthorized source.

RIP and RIP-2 in the network environment

Only one of the RIP versions should be used within a region. That is, a subnetwork can be configured to use either RIP-1 or RIP-2 but should not try to use both. The version of RIP in use should not be allowed on the subnetwork as it could cause conflicts. This means that all router links within the same subnet must be configured to use the same RIP version.

Note Different links on a router can use different RIP versions. In this case, the router performs conversion between the two RIP versions.

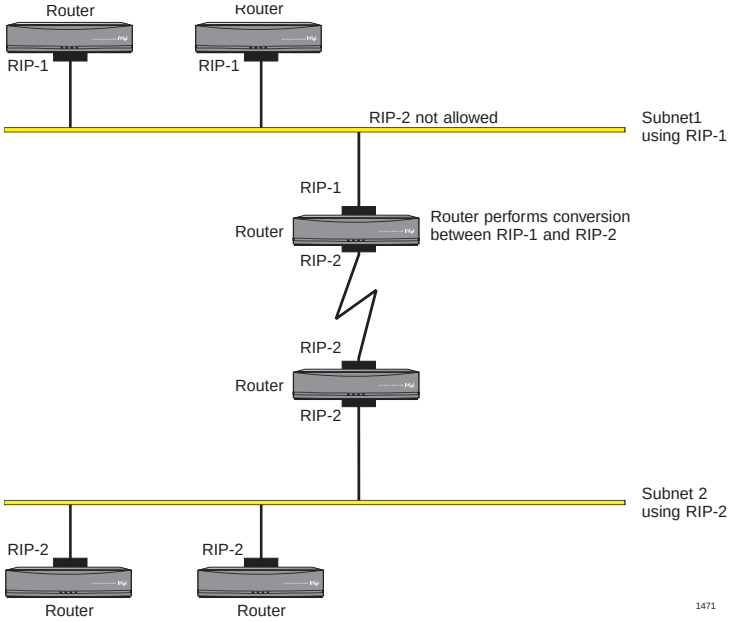

RIP implementation in the Intel Express 8100 Router

RIP-1 or RIP-2 is selected on an IP-link basis within the Intel Express 8100 Router during configuration. That is, RIP-1 or RIP-2 is selected for each router link configured for IP routing services during configuration.

RIP Metrics

RIP metric

Each router link has an associated RIP metric that may be set between the values 1 and 15. For LAN links the RIP metric is usually set to 1, for WAN links to a higher value but 1. Each RIP route has an associated metric value. A metric value higher than 15 means that the network is unreachable. The metric range limits the maximum diameter of the network.

Example of the correct use of metrics

The following illustration shows the correct use of metrics for route calculations. When calculating the best route using the metrics, data transmitted from Subnet 1 to Subnet 2 is transmitted over the 2 Mbps links, as the combined metric value path via the 2 Mbps links (metric = 1 for each link) is smaller than the metric value for path via the 64 kbps link (metric = 4).

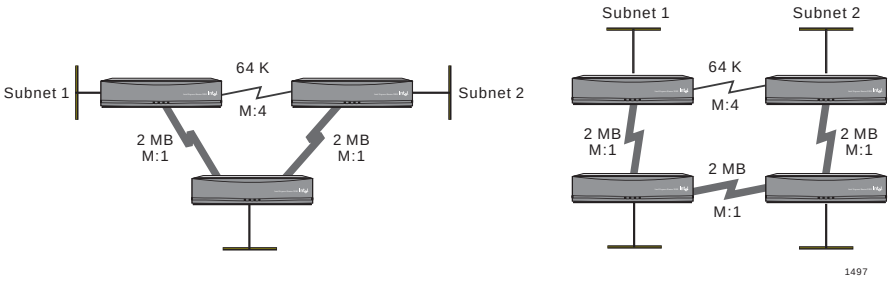

If the any of the 2 Mbps links are down, the 64 kbps link is used.

Example of metrics used incorrectly

The following example indicates how an inappropriate use of metrics may result in ineffective use of the available bandwidth. Here, data from Subnet 1 to Subnet 2 is transmitted across a 64 kbps link with metrics set to 3 while two (three) 2 Mbps links are available with their metrics set to 2:

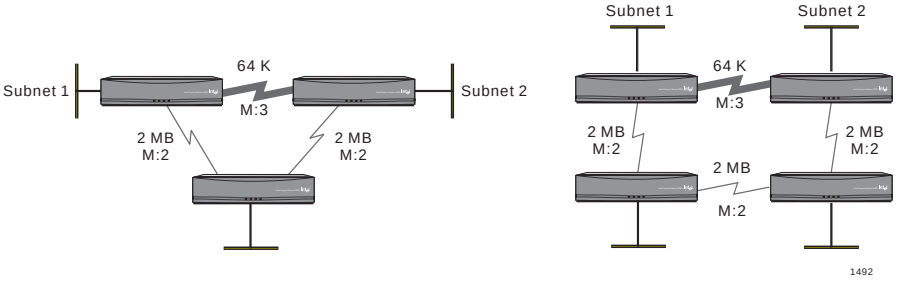

In this example, if the 64 kbps link is down, then the 2 Mbps links are used.

Triggered RIP

Introduction

RIP updates over WAN links should be minimized in order to minimize the cost of operation and maximize the available bandwidth for data communications. RIP updates can be an additional problem over dial-up (switched) WAN links as the link can often be brought up just to exchange routing information; the cost of operation can therefore be excessive.

Triggered RIP

Standard RIP updates are transmitted between routers at regular intervals (30 seconds) and whenever a network topology change is detected. Triggered RIP exchange routing information between routers whenever a WAN link is first brought up to ensure that the routing information tables are synchronized. After this initial synchronization, triggered RIP only exchanges routing information when a network topology change is detected. Triggered RIP therefore minimizes the use of WAN links for the exchange of routing information.

Multiple paths for triggered RIP

For devices using standard RIP only one path between networks is maintained in the routing table. If two paths to the same network are found by a device, the path with the lowest metric is stored or one of the paths is chosen if they have the same metric. When a device is using triggered RIP, all paths between networks are stored and secondary paths can be used if the primary path should fail. This is shown in the following example:

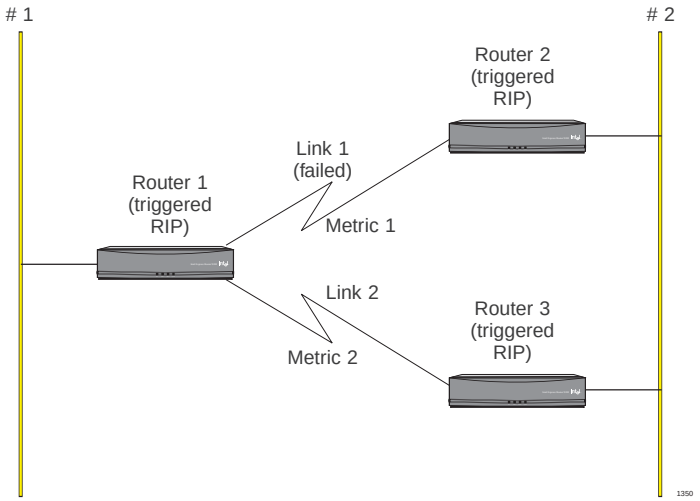

In the above diagram, router 1 must know both the paths to network 2. Otherwise, if Link 1 should fail, the router does not know about the path via link 2.

Static Routes

Definition

A static route is a permanent entry in the routing table.

Purpose

The purpose of setting a static route is to create a permanent route to an IP network or host. This can be done for the following reasons:

- to eliminate Routing Information Protocol (RIP) overheads from a link
  Eliminating RIP overheads from a link maximizes the available bandwidth for data transfers. This can be particularly useful for slow WAN links which could be swamped by the RIP updates. Using static routing to eliminate RIP is only practical for connecting to a few remote networks.
- to establish routes to IP networks and hosts which would not otherwise be reachable using RIP

  Two examples of where a static route must be established in order to reach destinations are given following.

IP RIP routing over PPP Tunnels

When using IP RIP routing over a PPP Tunnel a static host route to the IP address of the remote router must be established or the tunnel will not be stable.

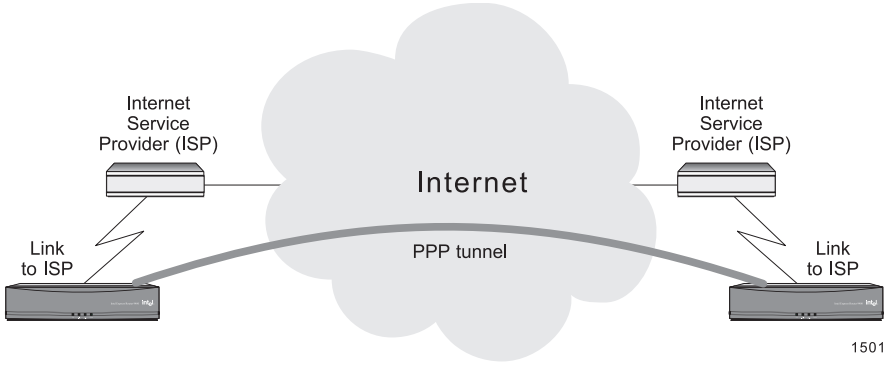

Remote boot provider

If a remote device (a device connected via a WAN link) is being booted via the router, a static route must be assigned to the network of the device being booted.

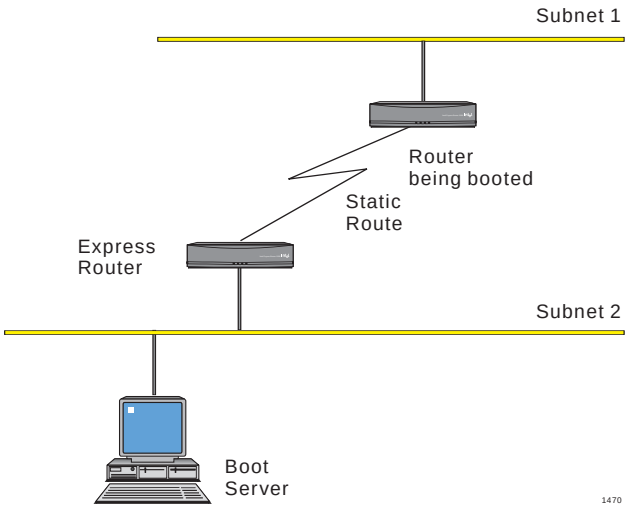

Characteristics

A static route is announced to other routers using RIP. A static route is floating so if another route with lower metric to the destination network is available, the route with the lowest metrics is taken.

Static routes and the RIP metrics

The metric used for static routes is assigned a default value by the router which can be redefined in the configuration program (valid input is from 1 to 15).

Static Routes Example

Goal and situation

A connection of subnet 2 and subnet 3 to the internetworks 17.0.0.0, 75.0.0.0 and 126.0.0.0 is desired. Only router X is connected to these internetworks. However, no connection to the internetworks 17.0.0.0, 75.0.0.0 and 126.0.0.0 via router X is possible at this point as router X does not run RIP.

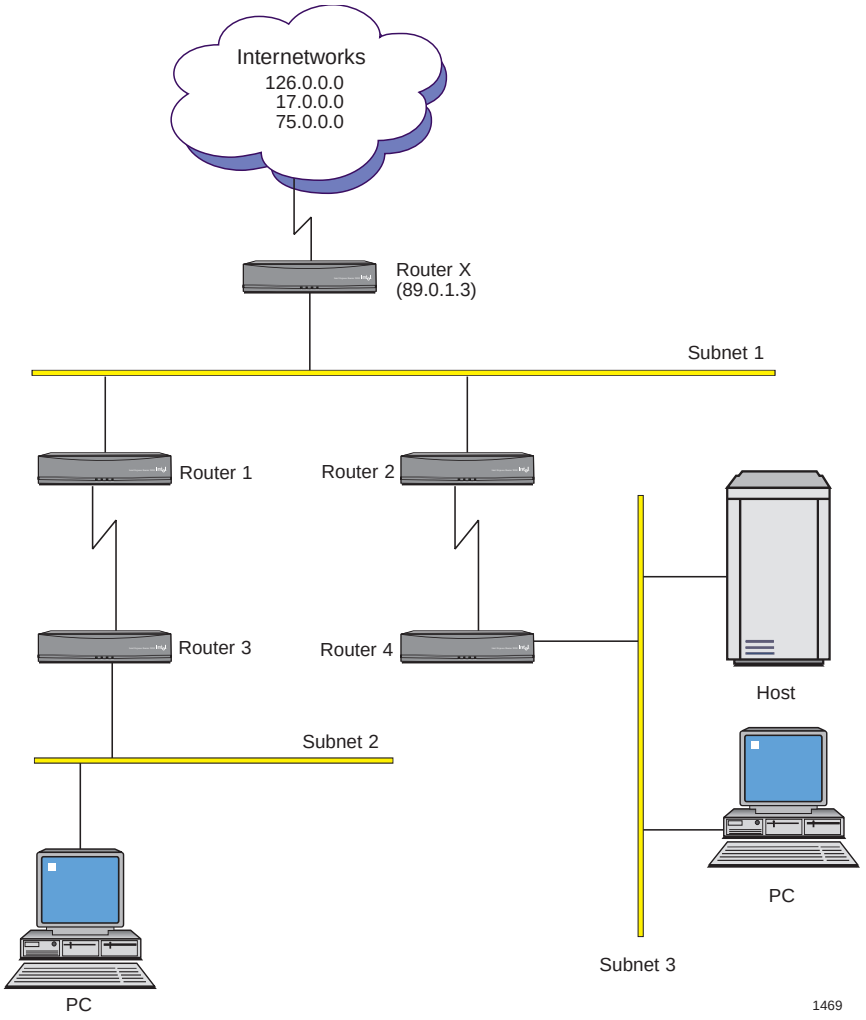

Solution

To establish a connection to the internetworks 17.0.0.0, 75.0.0.0 and 126.0.0.0 for subnets 2 and 3, a static route to the nets 17.0.0.0, 75.0.0.0, 126.0.0.0 via Router X must be created for both routers 1 and 2. Routers 1 and 2 then announce these static routes via RIP, meaning that the route to the internetworks 17.0.0.0, 75.0.0.0 and 126.0.0.0 is known by Routers 3 and 4 which make the internetworks reachable from subnets 2 and 3.

Chapter 9

Novell IPX Routing

In this chapter

This chapter describes Novell IPX/SPX Routing services in the Intel Express 8100 Router.

Novell Routing Concept

Concept description

Novell IPX routing is based upon a Routing Information Protocol (RIP) and Service Advertising Protocol (SAP).

Note Note that Novell IPX RIP is not the same as TCP/IP RIP.

Novell IPX calculates the best route to a destination based upon routing time delays associated with the links and with the number of hops (intermediate routers) taken to reach a network.

Background knowledge

It is assumed that you have a basic knowledge of IPX and the addressing system used in Novell IPX.

References

For further information on IPX, see Novell's "IPX Router Specification" Part Number: 107-000029-001.

Novell protocol overview

The following illustration provides an overview of Novell protocols:

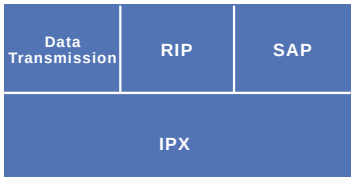

A short description of these protocols is given here:

Protocol	Description
IPX	The Internetwork Packet eXchange protocol is based on a network layer datagram without delivery guarantee. This protocol is used by all the protocols described above.
RIP	The Routing Information Protocol is used to gain access to remote networks and to exchange routing topology information.
SAP	The Service Advertising Protocol (SAP) is Novell's implementation of name-service. All services (for example, file servers, printer servers and gateways) announce themselves in periods of one minute.
Data Transmission.	For example NCP, SPX or NetBIOS.

Internetworking Packet Exchange (IPX)

Frame Types

Introduction

Novell supports different types of Ethernet frames. The frame type supported depends on the default value of the specific Novell driver.

The different frame type supported by Novell IPX are described in Appendix E "Novell IPX Frame Formats", p. 127.

To change the frame type

The frame type supported by your PCs can be changed using a program delivered together with the operating system for your NetWare server. These programs allow you to configure Ethernet IEEE 802.3, Ethernet IEEE 802.2, Ethernet SNAP and Ethernet II frames. The program available depends on the version of the Novell operating system.

Novell Addressing

Address notation

The Novell IPX hexadecimal address notation consists of a 4-byte network number, followed by the 6 byte node-ID (MAC address) and a 2-byte socket number, for example `00000001:000080123403:0451`, where `00000001` is the network number, `000080123403` the node ID and `0451` the socket number.

Network number assignment

A globally unique network number must be assigned to each network connection of a NetWare* server. Servers and router links attached to the same physical network must be configured with the same network number. To simplify network management, only servers and routers must be configured with a network number, as workstations automatically get their network number assigned by the server (connection to local server) or router (connection to remote server).

Sockets reserved by Novell

Each service reserved is allocated to a specific socket number. The following sockets are reserved by Novell:

Socket Numbers	Description
File Servers	
451h	NetWare Core Protocol (NCP) Process
Routers	
452h	Service Advertising Protocol (SAP) Process
453h	Routing Information Protocol (RIP) Process
Workstations	
455h	Novell NetBIOS Process
456h	Diagnostics Process
4000h to 7FFFh	Dynamic assigned sockets used by workstations for interaction with file servers and other network equipment.
8000h to FFFFh	Well-known sockets assigned by Novell.

Novell Metrics

Delays and hops

The Novell metrics are used to calculate the optimal route between networks (for IPX RIP) and services (for IPX SAP). Novell metrics use the following to calculate the best routes:

Delays (Ticks)

The delay is the time associated with sending data via a route, that is, if alternative routes are available the data packet is transmitted via the route that takes the least time.
The delay is measured in time ticks where 1 second = 18.21 time ticks.

Hops

The hops is the number of intermediate networks that must be passed in order to reach the destination network. Hops are only used in the unlikely event that two routes have an identical time tick; the route with the least number of hops is then be used.

The hops are also used to discard packets in case of looping, see “Looping packets” in “Data Transmission”, p. 76.

Metrics for RIP and SAP

Separate metrics (tics and hops) are calculated for dynamic routes (established using IPX RIP) and dynamic services (established using IPX SAP). These metrics are calculated automatically by the RIP and SAP protocols.

Metrics for static IPX routes and services

For static IPX routes and services, the time ticks and hops to networks and services cannot be calculated by the router. These parameters must therefore be manually calculated and defined for each static route and service.

Manually calculating metrics for static routes and services

When configuring static IPX routes and services, the best way to obtain the metrics for static routes and services is by enabling dynamic routing (using RIP and SAP) and reading the metrics from the routing and service table for the router, see “Establishing static routes and static services” in “Static Routes and Services”, p. 90. However, in some cases the metrics for a link may need to be calculated manually.

Note The delay for a route or service is the accumulated delays for all the paths to the route or service.

The metrics can be calculated manually as follows:

Delays (Ticks)

The delay for a link is associated with the speed (baud rate) of the link and can simply be calculated as:
`Link Delay (ticks) = 1100 / Link Speed (kbps).`
For example, a 64 kbps WAN link would have the following delay:
`1100 / 64 = 17.2` ticks.
This is only an approximation of the link delay calculated by the router (the router calculation is more complex and would assign a delay of 18). You should therefore try to use the delay calculated by the router (for the dynamic route or service you are trying to replace) wherever possible.

Hops

The hops is simply the number of intermediate networks that must be passed in order to reach the destination network.

Example

In the Novell metrics example, a data packet is sent from net 1 to net 4 via net 2 and 3, as the Delay (transmission time) is less (transmission speed on Ethernet 10 Mbps) than over the link with a baud rate of 9600 bps.

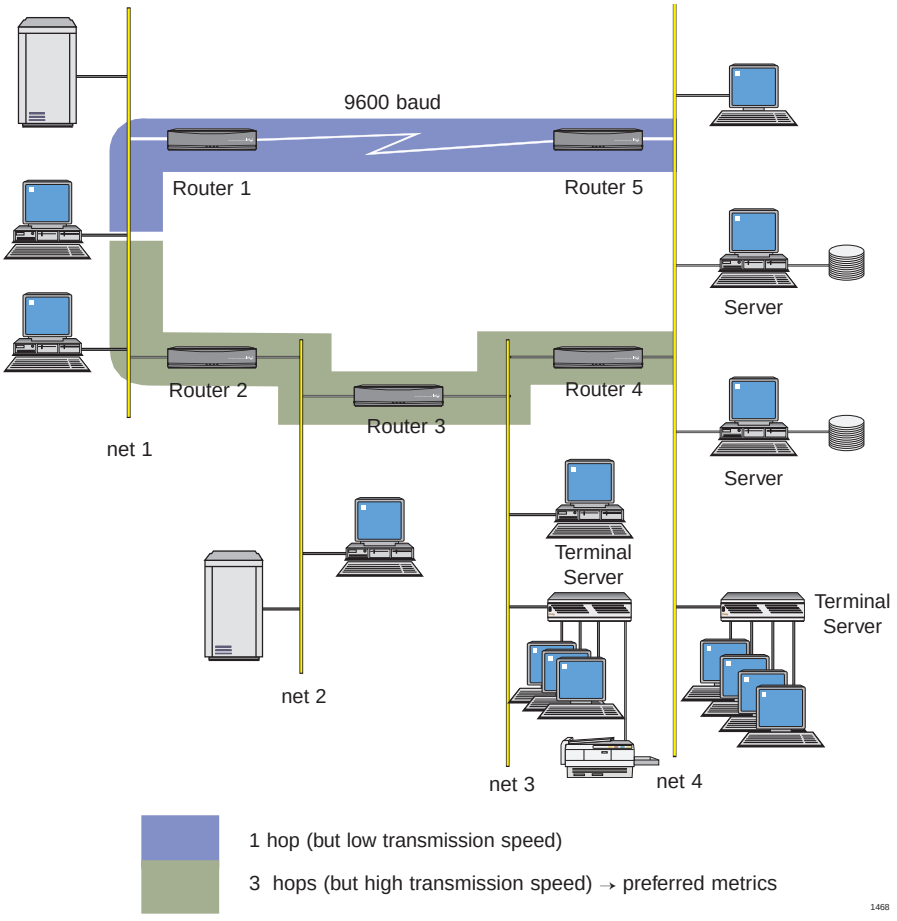

Data Transmission

Data transmission process

To send data from a NetWare device to a destination address on an other network, the following process is initiated:

1. The NetWare device (source address) broadcasts a routing request to all routers on the Local Area Network (LAN).
2. All local routers with a path to the destination address send a routing response to the querying workstation.
3. The NetWare device selects the best path and sends the data to the selected router.

Looping packets

If a packet is looping due to changes in the network topology, the packet is discarded on the 16th hop, as the hop count is incremented by every router.

IPX over WAN Links

Numbered and unnumbered WAN links

The Intel Express 8100 Router allows IPX WAN links to be either unnumbered or numbered where an IPX Network Number is assigned to the link. If the WAN Protocol is not being used, the routers over a WAN link must both be configured to run the same link type (unnumbered, numbered or unnumbered static). The negotiation of whether to use unnumbered or numbered links by the WAN Protocol is described in the following section.

Recommendations for the use of numbered links

If communicating with a router that supports unnumbered links (including all Intel router products), an unnumbered link should be used to avoid having to assign a network number to WAN links. If communicating with a router which only supports numbered links, a numbered link should be used.

IPX WAN Protocol

The IPX WAN Protocol can be selected for communicating over WAN links to ensure consistency with the configuration of the router on the other end of the WAN link. If the IPX WAN Protocol is selected, then it must be configured in both the routers over the WAN link—see the following section.

IPX WAN Protocol

IPX WAN 2 Protocol

The Intel Express 8100 Router supports the IPX WAN 2 protocol as defined by Novell in Novell NetWare Link Services Protocol Specification Rev 1.0, Chapter 3—IPX WAN Version 2, Novell part No. 100-001708-002. The protocol is backwards compatible with RFC 1634.

The protocol helps ensure consistency between the configuration of the two routers over the WAN link.

Numbered and unnumbered WAN links

The IPX WAN Protocol implemented in the Intel Express 8100 Router supports both unnumbered and numbered links. When a numbered link is defined, an IPX Network Number is assigned to the link. When protocol negotiation is taking place between the routers over a WAN link, unnumbered links are chosen in preference to numbered. Numbered links are only chosen if one or both of the routers cannot support unnumbered links.

Negotiation over a WAN link

When the WAN link is initialized, the IPX WAN protocol running in the routers at either end of the WAN Link negotiates various parameters to ensure that the routers are communicating on the same terms. Negotiation proceeds as follows:

1. Role Determination
   During this part of the negotiation, one router is designated as the ‘Slave’ and the other becomes the ‘Master’. If only one of the routers is configured for a numbered WAN Link, this router becomes Master. If both routers are configured for the same link type (numbered or unnumbered) the router with the highest Internal Net Number becomes the Master.
2. Parameter Negotiation
   When the roles of the routers have been determined, the Master attempts to find a protocol (numbered, unnumbered or static) which is acceptable to both the routers. If it succeeds, it informs the Slave of the protocol selected.
   If a numbered link is selected during negotiation, the Master also informs the Slave of the network number assigned to the link, ensuring that both routers use the same number for the link.
   The Master also informs the Slave of the delay used on the link.

IPX Watchdog Packets

Connections to a server

When a client is connected to a server, the session is registered in a table. It is a function of the server to keep this table up to date when client sessions are established and terminated.

Watchdog packets

When a client has not communicated with a server within a predefined time interval, the server sends out a message to the client to ask if it is still connected. This message is known as a watchdog packet. Watchdog packets can be sent by the IPX protocol (IPX watchdog packets) or the SPX protocol (SPX keep-alive packets—see “Sequenced Packet Exchange (SPX)”, p. 79).

If a client fails to respond to a predefined number of watchdog packets, the session is terminated at the server.

IPX watchdog packets over WAN links

A problem with IPX watchdog packets occurs when communication between clients and a server is over a dial-up (switched) WAN or ISDN link. If a number of clients are operating over the link, the link may never get the chance to become inactive and so the operating costs can be high. In addition, if the link is brought down for some other reason, the client sessions may be terminated prematurely.

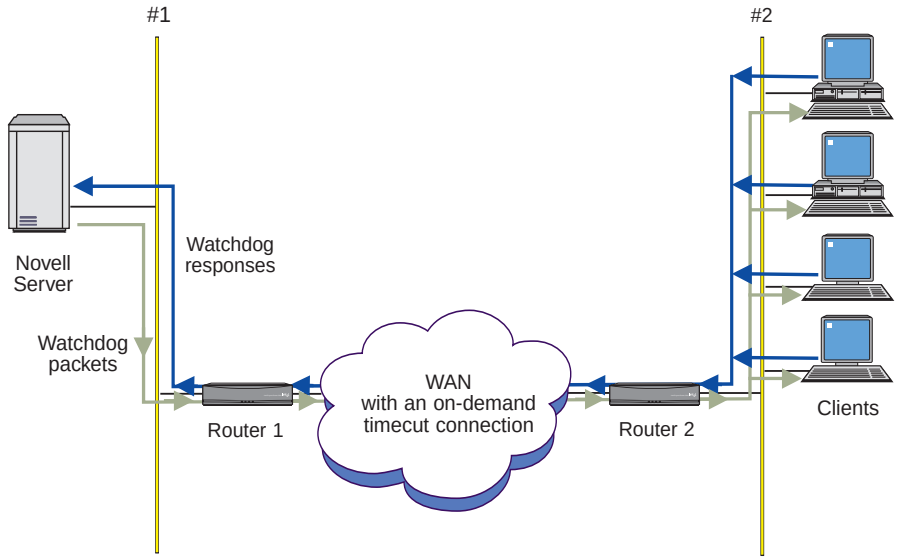

IPX watchdog spoofing

To counteract the problem of sending watchdog packets over dial-up (switched) WAN links, IPX watchdog spoofing can be set up in the router before the WAN link (on the server's side). When IPX watchdog spoofing is enabled, the router answers IPX watchdog packets on behalf of clients. The WAN link is therefore not activated to send IPX watchdog packets to clients.

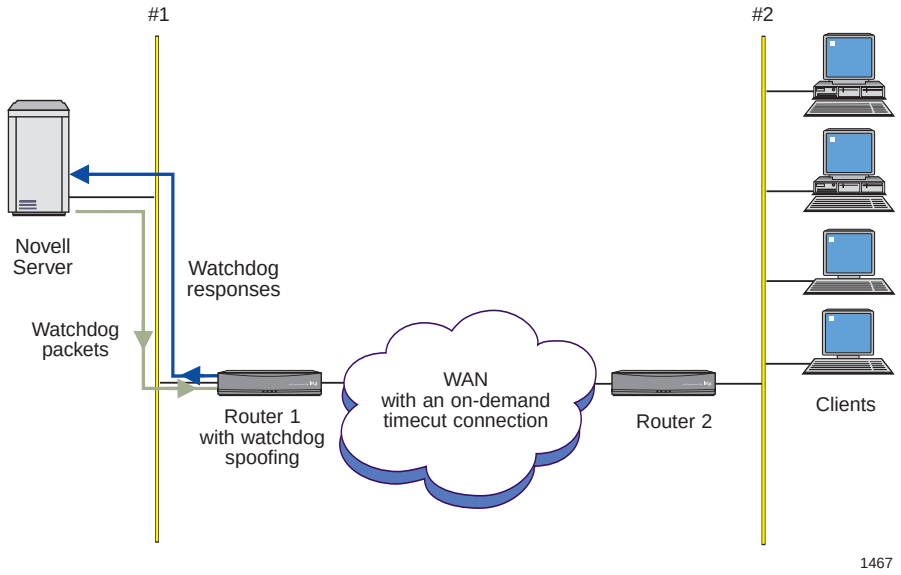

IPX Serialization Packets

IPX serialization packets

IPX Serialization Packets contain license information and are regularly sent across a network by NetWare servers (addressed to socket 457). If a server receives a Serialization Packet containing the license number it is using itself, it complains about license violations.

Discarded by the Intel Express 8100 Router

As IPX Serialization Packets are sent by NetWare servers at regular intervals, and do not allow dial-up (switched) WAN links to become inactive—the operating costs can therefore be excessive. For this reason, IPX Serialization Packets are discarded by the Intel Express 8100 Router.

Sequenced Packet Exchange (SPX)

Sequenced Packet eXchange (SPX)

SPX is a transport layer protocol which runs on top of the Internetwork Packet eXchange (IPX) protocol (network layer), and is used to control sessions between network devices across the network. An important function of the SPX protocol is guaranteed data delivery, as this function is not provided by the IPX protocol.

Sequenced Packet eXchange (SPX) is equivalent to the Transmission Control Protocol (TCP) in the IP protocol suite.

SPX Keep-alive packets

When using IPX/SPX protocol, the keep-alive function is taken care of by the SPX transport protocol, rather than by the IPX network protocol. As with IPX, SPX uses keep-alive packets to keep sessions alive—when an established session has not communicated within a predefined time interval, SPX keep-alive packets are used to determine if the session is still established. If an established session fails to respond to a predefined number of SPX keep-alive packets, the session is terminated. Unlike IPX watchdog packets, SPX keep-alive packets can be sent by all the network devices in the session.

SPX Keep-alive packets over WAN links

As with IPX watchdog packets, SPX keep-alive packets can cause problems when communicating over dial-up (switched) WAN links. If a number of SPX sessions are established over a WAN link, the link may never get the chance to become inactive and so the operating costs may be high. In addition, if the WAN link is brought down by some other reason, the SPX sessions may be terminated prematurely.

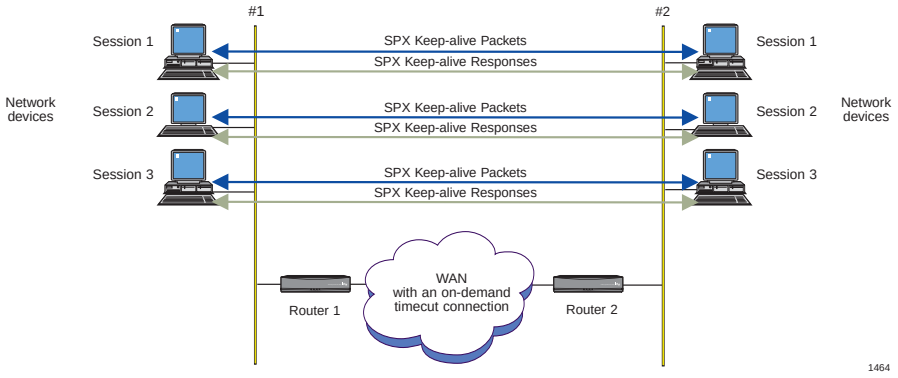

SPX spoofing

To counteract the problem of sending SPX keep-alive packets over dial-up (switched) WAN links, SPX spoofing can be set up in the routers over the WAN link.

Note SPX spoofing must be enabled in both routers over the WAN link for it to be effective.

When the WAN link is timecut and SPX spoofing is enabled, the router answers SPX keep-alive packets on behalf of devices over the WAN link. The WAN link is not activated just to send SPX keep-alive packets.

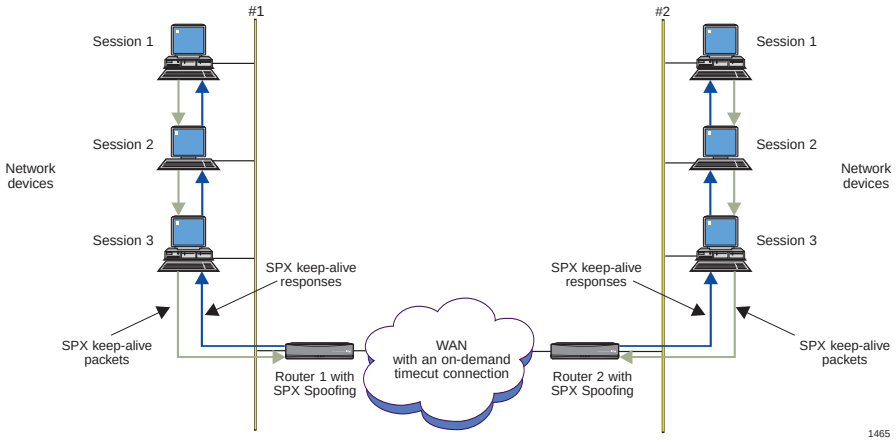

Terminating SPX sessions when SPX spoofing is enabled

When SPX spoofing is enabled, an SPX session could be maintained in the session tables of the network devices indefinitely. Most devices have a limited number of SPX sessions that can be established at any time, and redundant sessions kept alive by SPX spoofing must therefore be avoided. To counter this problem, a configurable SPX session timeout is initiated for each session which disables SPX spoofing after a defined time period. Network devices must then respond to SPX keep-alive packets as usual:

- If the SPX session is no longer valid it is terminated
- If the SPX session still valid, it responds to SPX keep-alive packets as usual and SPX spoofing is then resumed (the SPX spoofing timer is restarted)

At midnight (0:00) SPX spoofing is disabled for all sessions to allow devices to answer SPX spoofing packets as normal. This is implemented as a fail-safe mechanism to terminate sessions which are not cleared by the SPX spoofing session timeout (for example because the SPX session timeout value is set too high).

IPX Data Filters

IPX data filters

IPX data filters are defined on a link basis in the Intel Express 8100 Router, and separate filters are implemented for transmit (for restricting IPX packets leaving the router on a link), and receive (for restricting IPX packets entering the router from a link).

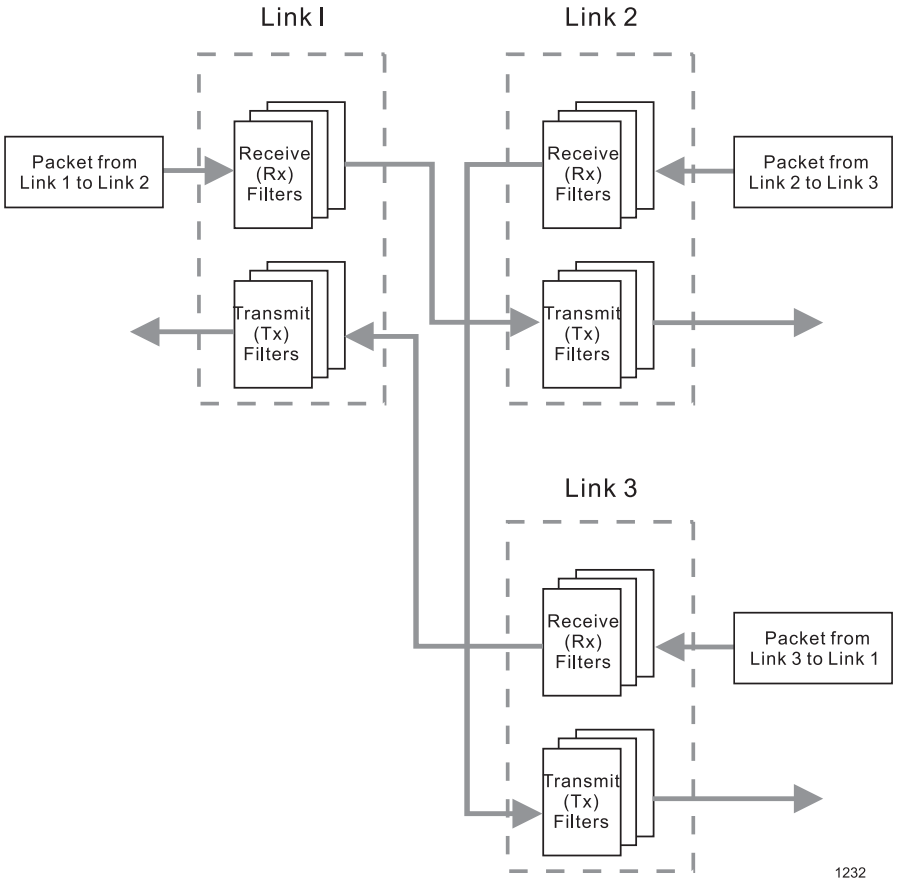

Note A good understanding of IPX is required before effective filtering can be set up on the router. Filters must be defined as strictly as possible to eliminate unauthorized access to services, hosts and networks.

Filter criteria

Each filter in the Intel Express 8100 Router can be setup to pass or discard IPX packets based on the following criteria:

Note All the filter criteria defined is used.

Destination Network Addresses

IPX data traffic addressed to a single destination network address, or range of network addresses specified by the `Destination Network Address` and `Destination Network Mask` parameters can be filtered.
When the `Destination Network Mask` parameter is set to FFFFFFFF, packets addressed to the single `Destination Network Address` specified will be filtered.
Filtering of packets to destination network addresses is disabled if a `Destination Network Address` of 00000000 with a `Destination Network Mask` of FFFFFFFF is defined (default values). For other values of the `Destination Network Mask`, data packets addressed to a range of network addresses will be filtered.
For example, with a `Network Address` of ABC00000 and a `Network Mask`: FFF00000, service information will be filtered from/to all servers on the networks beginning with ABC.

Destination Node Address

IPX packets addressed to a single device (node) can be filtered. A `Destination Node Address` of 000000000000 disables filtering on the destination node address.

Destination Socket

IPX packets addresses to specific sockets can be discarded. The most common IPX sockets (`NCP`, `NetBios`, `NLSP` and `IPX Ping`) can be selected from a list. Other sockets can be filtered by selecting `Other` for the `Destination Socket` and then entering the socket number `Value`.

Source Network Addresses

IPX data traffic from a single source network address, or range of network addresses specified by the `Source Network Address` and `Source Network Mask` parameters can be filtered.
When the `Source Network Mask` parameter is set to FFFFFFFF, packets from the single `Source Network Address` specified will be filtered.
Filtering of packets from source network addresses is disabled if a `Source Network Address` of 00000000 with a `Source Network Mask` of FFFFFFFF is defined (default values). For other values of the `Source Network Mask`, data packets from a range of network addresses will be filtered.
For example, with a `Network Address` of ABC00000 and a `Network Mask`: FFF00000, service information will be filtered from/to all servers on the networks beginning with ABC.

Source Node Address

IPX packets from a single device (node) can be filtered. A `Source Node Address` of 000000000000 disables filtering on the source node address.

Source Socket

IPX packets from a specific source socket can be discarded. The most common IPX sockets (`NCP`, `NetBios`, `NLSP` and `IPX Ping`) can be selected from a list. Other sockets can be filtered by selecting `Other` for the `Source Socket` and then entering the socket number `Value`.

Packet Type

Specific IPX packet types can be filtered. NCP and SPX packets can be selected from a list. Other packets types can be filtered by selecting `Other` for the `Packet Type` and then entering the socket number `Value`.

IPX Filters Example

Filter requirement

NetWare devices use Diagnostic Request Packets to get information from about other reachable devices. These packets can be used for testing if a device is reachable and gathering configuration information about devices.

Diagnostic Request Packets can be broadcast to all devices and require each device to respond and can cause WAN links to become saturated. It is therefore sometimes desirable to eliminate these diagnostics from WAN links, and this can be done using the IPX data filters implemented in the Intel Express 8100 Router.

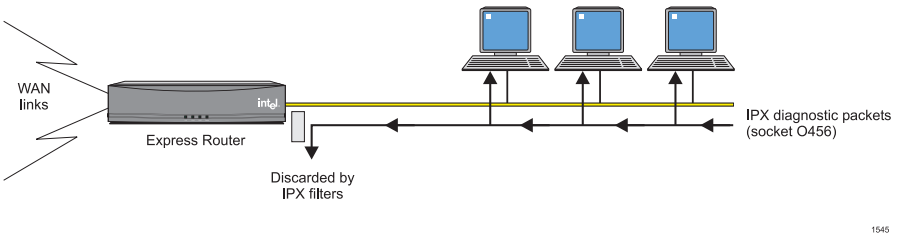

Note IPX ping may be affected by these filters, depending on whether the implementation of ping in the device which is testing for connectivity, uses the diagnostics responder.

Filter Required

To discard diagnostic packets from the LAN so that they are not transmitted on WAN links, set up a filter on the LAN link as follows:

- Set the default `Action` for the IPX Rx (receive) LAN filters to `Pass` to allow all data traffic to enter the router unless discarded by a filter.
- Set up an IPX Rx Filter for the IPX LAN link with the following criteria:
  - Set the filter `Action` to `Discard`
  - Set the `Destination Socket Value` to `0456` (the NetWare socket number for diagnostic packets)

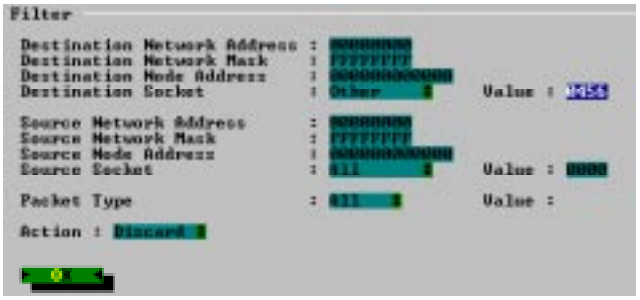

Routing Information Protocol (RIP)

Routing Information Protocol (RIP)

Routing information between routers is exchanged by means of the Routing Information Protocol (RIP) initiating the following functions:

- locates the fastest route
- retrieves routing information from other routers.
- Responds to requesting routers
- informs about the latest internetwork configuration
- informs about any topology changes in the internetwork

Encapsulation

The Intel Express 8100 Router uses ISO-9577 (SNAP) encapsulation when communicating over WAN Links. When talking to other routers via a WAN Link, they must also use (or be configured for) ISO-9577 (SNAP) encapsulation.

Routing table

A routing table contains a map of the entire network topology and is the tool for routers to provide each other with routing information.

Routing Information

Routers exchange routing information to inform each other about the network topology in the following situations:

- When powered on, initial broadcasts to directly attached segments and WAN links are sent.
- To receive routing information from other routers, Initial Routing Requests are sent.
- To inform about each other's existence, periodic broadcasts are typically sent every minute.
- Immediately after a network topology change, triggered updates are sent.

Updating periods

Each router sends updated routing information in periods of one minute (by default) on its associated LAN. In case of a router shut-down, all other routers delete the routing entry after approximately 3 minutes (or 3 times the normal interval of 1 minute). This allows a router to miss at least 2 packets in case the network is heavily loaded. Note that routers never send routing information back to networks they received the routing information from, as this would create loops.

Triggered updates

A triggered update is routing information exchanged due to changes in the network topology, for example if a link fails.

Cost factor

The costs for exchanging routing information on packet switched networks can be high, as the carrier charges you for every byte transmitted.

RIP Filtering

Introduction

Routing information entering and leaving a Intel Express 8100 Router can be filtered to remove unwanted routing information. For example, this may be done to:

- restrict access between certain networks
- reduce the amount of routing information traffic

Implementation

RIP filters are implemented for both WAN and LAN links, and are implemented for both received (Rx) and transmitted (Tx) routing information.

RIP Rx filters

RIP Rx filters can be used to remove selected incoming routing information before it is entered into the routing table for the router. Routing destinations removed in this way are not accessible via the router, and the routing information is not passed on to other routers.

RIP Tx filters

RIP Tx filters can be used to remove selected routing information before it is passed on to other routers. Routing destinations removed in this way are accessible via the router but the routing information is not passed on to other routers.

Service Advertising Protocol (SAP)

SAP functions

The Service Advertising Protocol (SAP) makes use of IPX and the medium-access protocols for its transport. SAP provides the following functions:

- initiates a workstation requests to get information about the name and address of the nearest server of a certain kind
- initiates a router requests to get information about the names and addresses of either all servers, or all servers of a certain kind on the internetwork
- sends responses to requests originating either from a router or a workstation
- initiates periodic broadcasts by servers and routers
- updates the accessible server information

SAP announces services and addresses of Novell NetWare networks. These service advertisements are collected by a SAP agent in a Server Information Table.

Service and Object types

A table indicates different kinds of services and associated objects that are currently known is given in Appendix F “Novell IPX Service Types”, p. 130.

How it works

Routers collect and exchange service information about all the services available on the local network using SAP agents. Workstations that require information about services available, send a broadcast, for example “get nearest server”. The local SAP agent then informs them about the nearest file server available.

SAP broadcasts are local broadcasts initiated by servers once a minute and received by local SAP agents only. That is, the packets are not forwarded beyond the local segments.

Service advertising

Novell services are offered from devices such as file servers (file transfer), printer servers (print service) and gateways (protocol conversion). Services are announced on the network. This service advertisement is not used directly by the end systems, but instead collected by SAP agents located in file servers and routers. SAP agents exchange information with the same method as the Routing Information Protocol and are therefore well informed about all active NetWare servers on the internetwork. A SAP agent is part of an IPX router/server.

SAP Filtering

Introduction

Service information entering and leaving a Intel Express 8100 Router can be filtered to remove unwanted services. For example, this may be done to:

- restrict access to servers from certain networks
- reduce the amount of service information traffic

Implementation

SAP filters are implemented for both WAN and LAN router links, and are implemented for both received (Rx) and transmitted (Tx) service access information.

SAP Rx Filters

SAP Rx Filters can be used to remove selected incoming service access information before it is entered into the service information table for the router. Services removed in this way are not be accessible via the router, and the service information is not passed on to other routers.

SAP Tx Filters

SAP Tx Filters can be used to remove selected service access information before it is passed on to other routers. Services removed in this way are accessible via the router but the service information is not passed on to other routers.

Static Routes and Services

Static routes and services

A static route is a route to a destination network which is always known to the router and does not need to be announced to the router using the Routing Information Protocol (RIP). Similarly, a static service is a service which is always known to the router and does not need to be announced to the router using the Service Access Protocol (SAP). Static routes and services are still announced to other routers using RIP and SAP as usual.

Use of static routes and services over WAN links

By establishing static routes and services over a WAN link and disabling RIP and SAP, no routing or service information is sent on the WAN link. This maximizes the available bandwidth for data communications, and for many WAN link types it reduces the costs of operation. Static routes are often used to prioritise the uses of WAN links for routing.

Note If RIP and/or SAP is disabled over a WAN link, static routes and services to all the networks and services over the WAN link must be established. The router may not be able to discover them otherwise.

When static routes and services are established, dynamic routing (using RIP and SAP) is usually disabled. An exception to this is to provide remote access to a network or service which is not otherwise announced to the router.

Static routes and static backup routes

The router can be set up to prioritise the use of static routes as follows:

- Static Routes (always use static route)
  The router always uses the static route (unless the static route is not available), even if a route with a lower metric is known. Routes to the same destination as the static route, announced to the router via RIP, are ignored unless the static route is down. If a number of static routes are available to the same network with the same priority, the router uses the static route with the lowest metric.
  When the static route is in use by the router, it is announced to other routers using RIP.
- Static Backup Routes (use best route)

  The router only uses a static backup route if a static route (primary static route) is not available and another route with a lower metric is not available. A static backup route is only announced to other routers using RIP if it is in use by the router (when it is in the routing table for the router).

Example use of a Static Backup Route

A typical example of the use of static backup routes is shown in the following diagram of access and general backup access between network 1 and 2. In addition, the PC on network 2 requires access and backup access to the server on network 1:

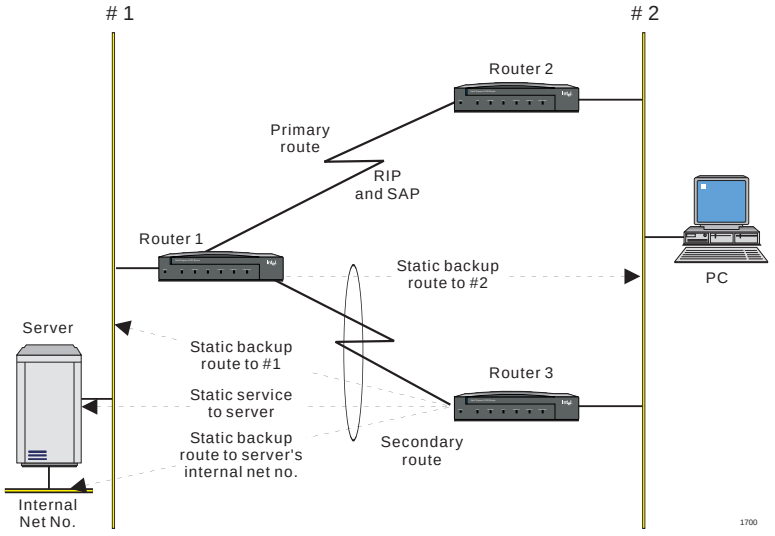

Example description

The above example is configured as follows:

- The primary link from network 1 to network 2 is the main route between network 1 and 2. This link uses dynamic routing (RIP and SAP).
- The secondary link between router 1 and router 3 is used as a static backup route between network 1 and 2 which is only used if the primary route fails. To allow general access between networks 1 and 2 in case the primary link should fail:
  - a static backup route to network 1 is configured in router 3
  - a static backup route to network 2 is configured in router 1
- To allow the PC to access the server on network 1, an additional static backup route to the internal network number of the server and a static service to the server must be established in router 3.

  Novell IPX servers use an internal network number which must be used when routing information to the server.
- For correct operation, the metrics for the backup routes must be higher than those for the primary route. To establish the static backup routes (in routers 1 and 3), dynamic routing over the backup link must be disabled and the static backup routes should be configured to `Use best route`. These routes are then only used if the primary route fails.

Dynamic routes with static backup routes consideration

For correct operation, a static backup route must have a higher metric than the primary route. If the primary route is a dynamic route using RIP, the metrics for the route are calculated automatically by the router. If the speed of the primary and backup links are similar, in certain circumstances the router can calculate the metrics for the primary route to be higher than the those for static backup link. If this occurs, the router forwards data via the static backup link rather than the dynamic link. To avoid this, the metrics for the static backup link should be set higher than the calculated value (for example, by about 50%).

Establishing static routes and static services

The easiest way to establish static routes, static backup routes and static services is to replace dynamic routes and services. This ensures that the static route/service information is correct (especially important for static services which can be quite complex), and that the metrics for the route/service are automatically calculated.

1. Configure the router to use dynamic routing (using RIP/SAP) over the link. This is done by setting the `Routing Protocol` parameter to `RIP/SAP`.
2. Obtain the relevant information for the static routes and services from the routing table for the router.
   For static routes (and static backup routes), the information required includes the network numbers and the associated metrics (delays and hops). For static services, the information required includes, the server names, types, network, node, socket and the hops to the server.

3 Reconfigure the router, disable dynamic routing on the link (if required) and configure the required static routes, static backup routes and services using the information obtained in Step 2.
Disabling dynamic routing on the link.

Configuration Example—Cost-Reduced WAN

Optimal cost reduced WAN example

The optimal use of the WAN link is achieved by disabling RIP and SAP over the WAN link and establishing static routes and services over the WAN link. This eliminates all routing traffic from the WAN link, optimising the available bandwidth for data transfers. However, establishing static routes and services is only practical if relatively few networks and services require accessing over the WAN link. For larger networks, dynamic routing using RIP and SAP is required to keep track of network topology changes.

Cost-reduced WAN using RIP and SAP example

An example of a cost reduced WAN example using RIP and SAP is described following.

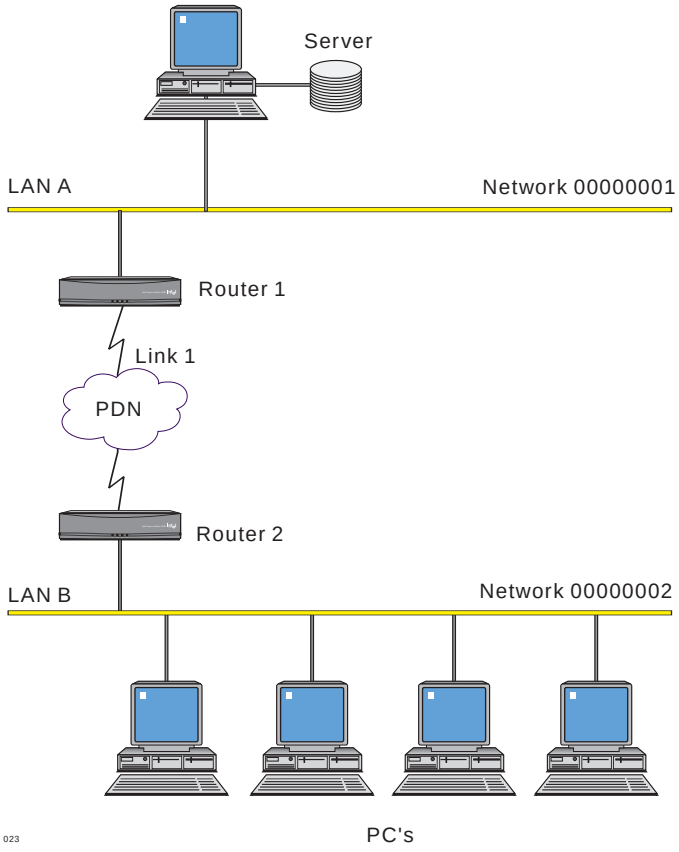

The parameter settings for a cost-reduced WAN using RIP and SAP for Router 1 and Router 2 is provided below. LAN configuration is a basic configuration.

Parameters	Router 1 -WAN	Router 2— WAN
Network number	unnumbered	unnumbered
RIP Triggered Updates	Disabled	Disabled
SAP Triggered Updates	Disabled	Disabled
RIP Update Interval	3600 seconds	3600 seconds
SAP Update Interval	3600 seconds	3600 seconds
Type 20 Broadcast	Disabled	Disabled
Watchdog Spoofing	Enabled	Enabled
Filters	see below	see below

When configuring a cost-reduced WAN connection using RIP and SAP, `Triggered Updates` (for RIP and SAP) should be disabled and `Update Intervals` (for RIP and SAP) should be set to a high value to reduce the use of bandwidth. `Watchdog Spoofing` and `SPX Spoofing` should also be enabled so that dial-up (switched) WAN links are be initiated just to send IPX watchdog packets and SPX keep-alive packets to remote clients.

Cost-reductions using RIP and SAP filtering

This cost-reduced WAN example can be further improved by using RIP and SAP filtering to reduce the routing and service access information exchange over the WAN link (see “RIP Filtering”, p. 86 and “SAP Filtering”, p. 87).

Chapter 10

WAN Bridging Services

In this chapter

This chapter describes configuring the Intel Express 8100 Router to use bridging services and to use Spanning Tree Algorithm.

Introduction to WAN Bridging

WAN Bridging

The Intel Express 8100 Router can bridge packets between the LAN and a WAN link or between the two WAN links—this is known as WAN Bridging.

Background knowledge

It is assumed that you have a basic knowledge of the terminology used in bridged environments.

Note When the router is operating as a bridge, it is referred to as a bridge in this chapter.

The purpose of bridging in the router

The purpose of WAN Bridging is to extend the router with bridging services so that the router can also act as a bridge.

Bridging concept

Bridging in the router conforms to the following standards:

- IEEE 802.1D—Local MAC Bridging
- IEEE 802.1G—Remote MAC Bridging

Bridging relies on MAC layer addressing (specified by IEEE 802.3) and is independent of Network layer protocols.
Bridging is used by the Intel Express 8100 Router for protocols which cannot be routed by the router. However, if a routing protocol is disabled, the protocol is bridged if bridging services are enabled.

Spanning Tree Algorithm

The Spanning Tree Algorithm is implemented for bridging in the router. The Spanning Tree Algorithm gives the possibility of having redundant paths within a bridged environment, allowing stand-by bridges to be implemented which are ready to start forwarding packets if another bridge is powered off or fails during operation.

The Spanning Tree Algorithm is described in "Spanning Tree Algorithm", p. 108.

Bridge Spoofing

Within bridged networks which implement the Spanning Tree Algorithm, configuration BPDU packets are transmitted between the bridges to communicate topology and other changes. For financial reasons, it is desirable to avoid the unnecessary transmission of configuration BPDU packets over WAN links.
A Bridge Spoofing feature is available in the Intel Express 8100 Router which stops the transmission of configuration BPDU packets over a WAN link unless they are important. This is known as BPDU Spoofing.

Another feature of Bridge Spoofing is that it stops the forwarding of Non-BPDU packets over WAN links to ports which are in a Blocking State. A Blocking State can be assigned to a bridges port by the Spanning Tree Algorithm and the bridge port then discards all packets it receives. Forwarding packets to a bridge port which is Blocking is therefore unnecessary and is a waste of WAN resources. This is known as Non-BPDU Spoofing.

Bridge Spoofing is described in "Bridge Spoofing", p. 114.

Loop-Back check

A loop-back check is implemented in this router to ensure that packets transmitted received back on the same link are discarded. This feature stops loop-backs in the network segments, for example, if a WAN link is setup to loop back packets.

How Bridging Works

When are packets routed and when are they bridged?

If a packet is received for forwarding by the router and the appropriate routing protocol is enabled, the packet is routed. Otherwise the packet is bridged.

Bridging TCP/IP and Novell IPX

The Intel Express 8100 Router is normally set up to route TCP/IP and Novell IPX. If routing of these protocols is disabled, the router automatically bridges them if bridging services are enabled. The router cannot route a protocol to between two links and bridge them between others, a routing protocol must be disabled completely before the protocol is bridged. However, IPX routing services must be configured for one particular Ethernet packet type (SNAP, TYPE, 802.2 or 802.3). All other IPX Ethernet packets types are bridged if bridging services are enabled.

Bridging and routing services

Bridging Services operate at the MAC Layer (also known as Data Link Layer or Layer 2). Bridging Services operate independently of Routing Protocols which are Network layer (Layer 3) services. Bridging can therefore be used to forward packets in any Network Layer protocol.

Bridging Services use the MAC addresses assigned to network devices, where routing protocols usually require protocol-specific addresses to be assigned to devices. This makes bridged networks much simpler to implement than routed networks from the Network Managers point of view.

The use of Forwarding Tables

Bridging Services use Forwarding Tables to decide where packets should be bridged to.

- If a packet received for bridging appears in one or more of the forwarding tables, it is forwarded to one or more links or discarded as appropriate.
- packets received for bridging which do not appear in the forwarding tables are transmitted to all the network segments except the one from which it was received. Alternatively, the bridge can be set up to discard packets which do not appear in the forwarding tables.

A full description of bridging using the forwarding tables is given in "Operation of the Bridge in the Intel Express 8100 Router", p. 101. Forwarding tables are described in the "Forwarding Tables", p. 95.

Forwarding Tables

Forwarding tables

Forwarding Tables are used when deciding where incoming packets should be bridged to. The following forwarding tables are implemented in the router:

- Destination Forwarding Tables for both Unicast and Multicast packet destinations
- a Source Forwarding Table for defining destinations for packets from source MAC addresses
- a Type Forwarding Table for defining destinations for Ethernet packet types

Destination Forwarding Tables

There are two destination Forwarding Tables implemented in the router—a Unicast Destination Forwarding Table and a Multicast Destination Forwarding Table. These tables consist of the network access rights for associated destination MAC addresses—Unicast addresses in the Unicast Forwarding Table and Multicast addresses in the Multicast Forwarding Table.
Each entry in both the Unicast and Multicast Forwarding Tables consist of a destination MAC address and the allowed incoming link(s) and outgoing link. For example:

Unicast Destination Forwarding Table example:

Unicast Destination Forwarding Table (example)		
Destination MAC Address	**Incoming Link**	**Outgoing Link**
000080000001	All	WAN link 2
0000800000AA	All	None
0000800000BB	All	WAN link 1
0000800000BB	All	WAN link 2
and so on.		

Multicast Destination Forwarding Table example:

Multicast Destination Forwarding Table (example)		
Destination MAC Address	**Incoming Link**	**Outgoing Link**
010080000002	LAN link	WAN link 1
01008000000C	WAN link 1	WAN link 2
01008000000D	WAN link 1	LAN link
01008000000E	All	LAN link
01008000000F	All	None
010080111111	LAN link	WAN link 1
010080111111	LAN link	WAN link 2
and so on.		

Notes:

- For entries in the Unicast Destination Forwarding Table, the incoming link can only be `All`, that is, packets having the defined Unicast MAC destination address are allowed from any of the links.
- For entries in the Multicast Destination Forwarding Table, the incoming link

can only be `All` when the outgoing link is `None`.

- Static entries are defined together for the Unicast and Multicast Destination Forwarding Tables and not in the separate tables.

Default bridge action for unknown destinations

The bridge can be set up to either forward to all links (default) or discard packets received for bridging which do not appear in the Unicast or Multicast Destination Tables, see "Controlled Bridging", p. 106.

Static and learned entries

Some of the entries in the Unicast Destination Forwarding Table CAN BE static entries while all the entries in the Multicast Destination Forwarding Table ARE static entries. Static entries are entries defined during configuration of bridging services.

Unless bridge learning has been disabled, the Unicast Destination Forwarding Table may also contain "learned entries". Learned entries are entries that the router has found out about by listening to the connected network segments—the router records the source MAC address from the packets transmitted on the network segments. Learned entries are automatically put into the Unicast Destination Forwarding Table and allow the bridge to forward packets to the specific network segment towards the Destination MAC address. Bridge learning is described in detail in the standard IEEE 802.1D.

Up to 2048 static and learned Unicast entries can be stored in the Unicast Destination Forwarding Table, up to 124 of them can be static entries. Up to 62 static Multicast entries can be defined in the Multicast Destination Forwarding Table.

Note If the Unicast Destination Forwarding Table becomes full, the next learned entry overwrites the oldest existing learned entry in the table. In this case, bridging services still works but do not operate optimally.

Disabling bridge learning

Bridge learning can be disabled so that the Unicast Destination Forwarding Table only contains Static entries defined during configuration of bridging services. In addition, the bridge can be set up to discard packets received for bridging which do not appear in the Unicast or Multicast Destination Tables. The bridge then only forwards packets to user-defined destination MAC addresses. This may be done, for example, for security purposes (to eliminate unauthorized traffic) and for restricting unnecessary or non-essential traffic.

Default static Multicast destinations

A number of factory default entries are included in the Multicast Destination Forwarding Table. These default entries are set to reject certain packets and packets used by routing and bridging protocols and are:. The default entries can be removed during configuration of the router.

Destination MAC Address	Explanation
010080000000	Proprietary Multicast messages.
010080000001	Proprietary Multicast messages.
010080000004	Proprietary Multicast messages.
0180C2000000	Spanning Tree Algorithm configuration BPDU packets.
0180C2000001	IEEE 802.1D demands that they are rejected.
0180C2000002	IEEE 802.1D demands that they are rejected.
0180C2000003	IEEE 802.1D demands that they are rejected.
0180C2000004	IEEE 802.1D demands that they are rejected.
0180C2000005	IEEE 802.1D demands that they are rejected.
0180C2000006	IEEE 802.1D demands that they are rejected.
0180C2000007	IEEE 802.1D demands that they are rejected.
0180C2000008	IEEE 802.1D demands that they are rejected.
0180C2000009	IEEE 802.1D demands that they are rejected.
0180C200000A	IEEE 802.1D demands that they are rejected.
0180C200000B	IEEE 802.1D demands that they are rejected.
0180C200000C	IEEE 802.1D demands that they are rejected.
0180C200000D	IEEE 802.1D demands that they are rejected.
0180C200000E	IEEE 802.1D demands that they are rejected.
0180C200000F	IEEE 802.1D demands that they are rejected.
0180C2000014	OSI IS-IS level 1 routers.
0180C2000015	OSI IS-IS level 2 routers.
09002B000004	OSI ES-IS end systems.
09002B000005	OSI ES-IS intermediate systems.

Source Forwarding Table

The Source Forwarding Table contains the network access rights for associated source MAC addresses. Each entry in the table consist of a source MAC address, the incoming link and allowed outgoing link. For example:

Source Forwarding Table (example)		
Source MAC Address	**Incoming Link**	**Outgoing Link**
000080222222	LAN link	WAN link 1
0000800000CC	LAN link	WAN link 2
0000800000DD	WAN link 1	WAN link 2
0000800000EE	All	WAN link 1
0000800000FF	All	None
and so on.		

Note For entries in the Source Forwarding Table, the incoming link can only be `All` when the outgoing link is `None`.

The Source Forwarding Table can only contain static entries defined by the operator during configuration of the router. Up to 62 entries can be defined.

Default bridge action from unknown sources

The bridge can be set up to either forward to all links (default) or discard Multicast packets from source addresses which do not appear in the Source Forwarding Tables. This feature which prevents Multicast packets from unknown source MAC addresses from being forwarded over WAN links, and can be used when a central site router is communicating with a number of remote devices.

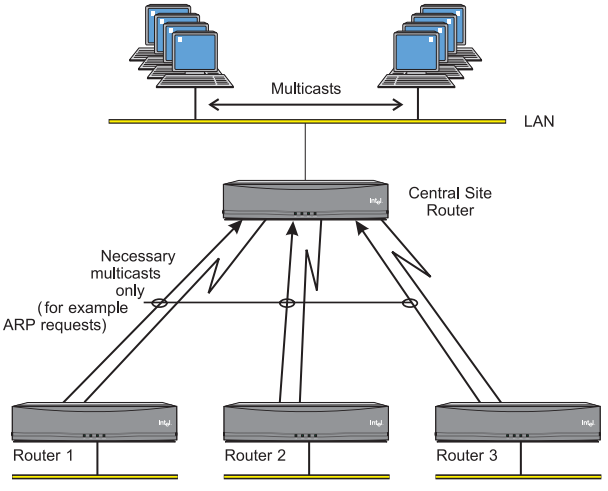

The Default Action in the central site router is set to `Discard Multicast`, and the MAC addresses of the remote routers are allowed to pass by filter entries. This allows necessary remote Multicast packets (using the MAC address of the router) to be forwarded by the router while all local Multicasts are discarded without being transmitted on the WAN links.

Typically, the default action of the Destination Forwarding Table is set to Discard to prevent stations that the router does not know about from being Multicast to all destinations.

Type Forwarding Tables

The Type Forwarding Table contains the network access rights for associated Ethernet packet types. Each entry in the table consist of an Ethernet packet type code the incoming link and the allowed outgoing link. The known Ethernet Type Codes are listed in Appendix C “Ethernet Type Codes”, p. 121. For example:

Type Forwarding Table (example)		
Ethernet packet Type	**Incoming Link**	**Outgoing Link**
1231	LAN link	WAN link 2
1234	WAN link 1	LAN link
1236	WAN link 1	LAN link
AAAA	WAN link 2	LAN link
BBBB	WAN link 1	WAN link 2
CCCC	All	None
DDDD	All	WAN link 1
and so on.		

Note For entries in the Type Forwarding Table, the incoming link can only be `All` when the outgoing link is `None`.

The Type Forwarding Table can only contain static entries defined by the operator during configuration of the router. Up to 62 entries can be defined.

Default bridge action for unknown packet types

The bridge can be set up to either forward to all links (default) or discard packets received for bridging which do not appear in the Type Forwarding Tables. The bridge then only forwards user-defined Ethernet packet types. This may be done, for example, for security purposes (to eliminate unauthorized traffic) and for restricting unnecessary or non-essential traffic.

Filtering Function

The static entries in the Forwarding Tables can be setup to discard packets, thereby acting as a filter. To do this, the outgoing link must be defined as `None` (see above example).

Forwarding tables used

Bridged packets are processed by all the Forwarding Tables before being forwarded. A packet which is allowed to pass by one Forwarding Table can therefore be rejected by another Forwarding Table.

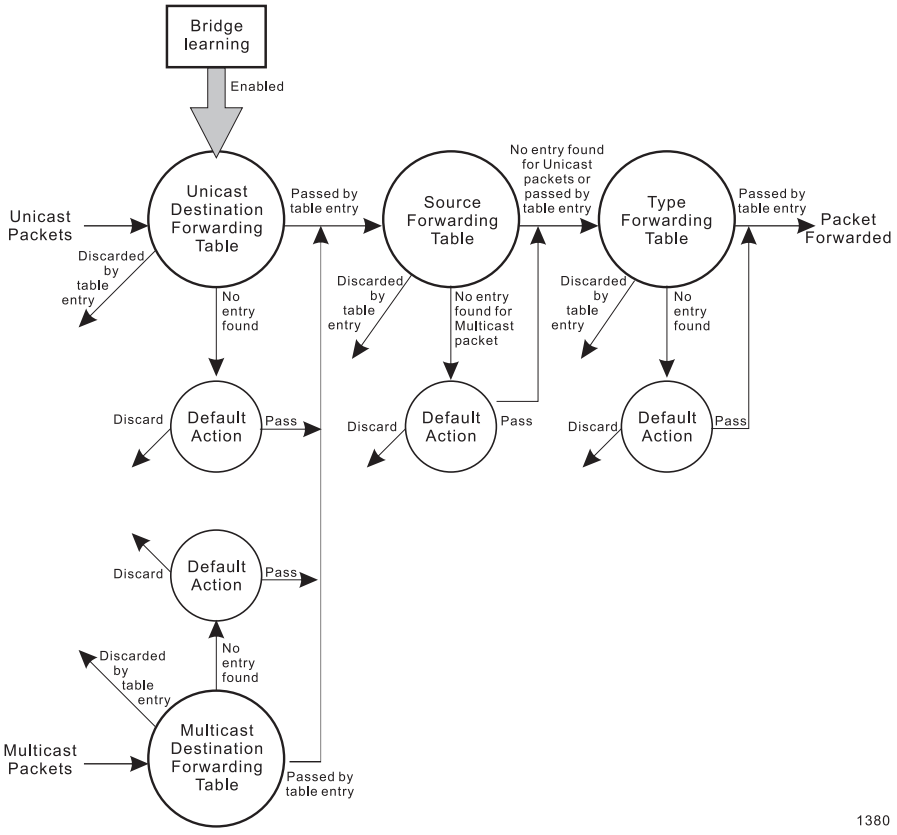

Operation of the Bridge in the Intel Express 8100 Router

Loop-Back check

When a packet is received for bridging, a check is performed to determine if the packet was bridged to the network segment by the router itself—this is called a “loop-back” check. Packets are discarded if they are received on the same link that they were transmitted.

Forwarding

A packet which is not discarded by the loop-back check is checked against the Forwarding Tables (as described in the previous section) to find out which network segment(s) it should be transmitted. The packet is forwarded on the link(s) allowed by the Forwarding Tables except the link from which it was received.

If the packet does not correspond to the entries in any of the Forwarding Tables, it can be forwarded on all the links except the link from which it was received or discarded, depending on the configured destination, source and packet type default bridging function.

Packet processing by the Destination Forwarding Tables

Packets which are received for bridging are checked against one of the Destination Forwarding Tables—packets with Unicast MAC addresses are only checked against the Unicast Destination Table and packets with Multicast MAC addresses are checked against the Multicast Destination Table. The process is as follows:

1. If table entries with both the destination MAC address and incoming link exist, the packet is labelled for forwarding on the outgoing link(s) defined by these entries. The packet is then passed to the Source Forwarding Table for further processing.

   Note If a table entry for the packet defines the outgoing link as `None`, the packet is discarded.

2. If a table entry exists for the Multicast destination MAC address but the incoming link does not correspond with that of the packet, the packet is discarded.

   Note Unicast destination MAC addresses are allowed from any link.

3. If no entry is found for the destination MAC address and the default action of the Unicast or Multicast Destination Table (as appropriate) is set to `Discard`, the packet is discarded.

4. If no entry is found for the destination MAC address and the default action of the Destination Table (Unicast or Multicast as appropriate) is set to pass, the packet is labelled for forwarding on all the links except the link on which it was received. The packet is then passed to the Source Forwarding Table for further processing.

   Note To avoid “storms”, a maximum of 5 packets with unknown Unicast destinations can be forwarded by the router within a 5 second period. Additional packets are discarded.

Processing by the Source Forwarding Table

Packets which are passed by the Destination Forwarding Table are then checked by the Source Forwarding Table as follows:

1. If a table entry for both the Source MAC address and the incoming link exist, the packet is labelled for forwarding on the outgoing link(s) allowed by both the Destination and Source Forwarding Tables. The packet is then passed to the Type Forwarding Table for further processing.
   For example, if an entry in the Destination Forwarding Table allows a packet from the LAN link to be forwarded on WAN link 1 and an entry in the Source Forwarding Table only allows the packet from the LAN link to be forwarded on WAN link 2, the packet received on the LAN link is discarded.

   Note If a table entry for the packet defines the outgoing link as `None`, the packet is discarded.

2. If a table entry exists for the Source MAC address but the incoming link does not correspond with that of the packet, the packet is discarded.
3. If Source Forwarding Table is set to `Discard Multicast`. If the packet is a Multicast packet and the Source MAC address for the packet is not found in the table, the packet is discarded.
4. If no entry is found for the Source MAC address and the Default Action of the Source Forwarding Table is set to Pass, the packet is labelled for forwarding on the link(s) defined during processing by the Destination Forwarding Table. The packet is then passed to the Type Forwarding Table for further processing.

Processing by the Type Forwarding Table

Packets which have been passed by the Destination and Source Forwarding Tables are checked against the entries in the Type Forwarding Table as follows:

1. If table entries for both the Ethernet packet type and the incoming link exist, the packet is forwarded on the outgoing link allowed by the Destination, Source and Type Forwarding Table entries except for the link on which it was received.

   Note If entries exist in the Destination, Source and Type Forwarding Tables, the outgoing link in these entries must be the same or the packet is discarded.

   Note If a table entry for the packet defines the outgoing link as `None`, the packet is discarded.

2. If a table entry exists for the Ethernet packet type but the incoming link does not correspond with that of the packet, the packet is discarded.
3. If no entry is found for the Ethernet packet type and the Default Action of the Type Forwarding Table is set to `Discard`, the packet is discarded.

4 If no entry is found corresponding to both the Ethernet packet type and the Default Action of the Type Forwarding Table is set to Pass, the packet is forwarded on the link(s) defined during processing by the Destination and Source Forwarding Tables.

Forwarding example 1

An incoming packet is received by the router for bridging on WAN link 1 The packet has a destination MAC address of 000080000001, a source MAC address of 000080000002 and a packet type code of 1234. It is processed as follows:

1. It is checked for loop-back. The packet has not been sent on the link on which it has just been received so it passes the loop-back test without being discarded.
2. The packet has a Unicast destination MAC address and so is checked against the entries in the Unicast Destination Forwarding Table (see “Destination Forwarding Tables” in “Forwarding Tables”, p. 96). An entry is found in the table for a destination MAC address of 000080000001. All incoming links are always used by the Unicast Destination Table. The entry defines the outgoing link as WAN link 2.
3. The source MAC address of the packet is checked against the entries in the Source Forwarding Table and is not found (see “Source Forwarding Table” in “Forwarding Tables”, p. 99).
4. The packet type code (1234) is checked against the entries in the Type Forwarding Table (see “Type Forwarding Tables” in “Forwarding Tables”, p. 100). An entry is found for packet type 1234 with an incoming link of WAN link 1. The entry defines the outgoing link as the LAN link.
5. The packet is discarded as the Unicast Destination Forwarding Table only allows the packet to be forwarded on WAN link 2 and the Type Forwarding Table only allows the packet on the LAN link. The packet is therefore incompatible with both these conditions at the same time.

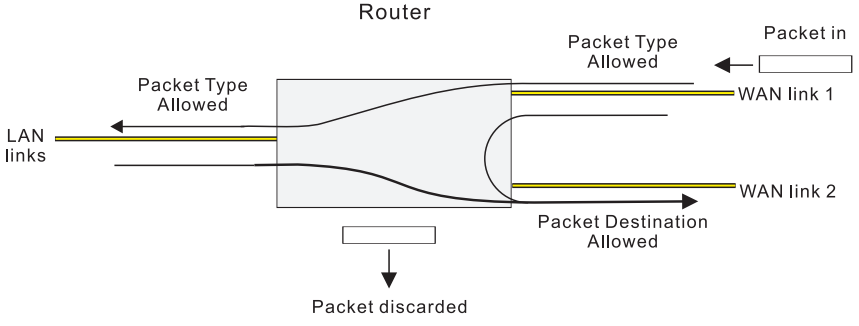

Forwarding example 2

An incoming packet is received by the router for bridging on the LAN link. The packet has a destination MAC address of 010080111111, a source MAC address of 000080222222 and a packet type code of 4321. It is processed as follows:

1. It is checked for loop-back. The packet has not been transmitted on the link on which it has just been received so it passes the loop-back test without being discarded.
2. The packet has a Multicast destination MAC address and so is checked against the entries in the Multicast Destination Forwarding Table (see "Destination Forwarding Tables" in "Forwarding Tables", p. 96). Two entries are found in the table for a destination MAC address of 010080111111 and the incoming link as the LAN link. One of the entries defines the outgoing link as WAN link 1, the other defines the outgoing link as WAN link 2. The packet is therefore labelled for forwarding on both the WAN links.
3. The source MAC address of the packet is checked against the entries in the Source Forwarding Table (see "Source Forwarding Table" in "Forwarding Tables", p. 99). An entry is found for a source MAC address of 000080222222 with the incoming link as the LAN link. The entry defines the outgoing link as WAN link 1.
4. The packet type code (4321) is checked against the entries in the Type Forwarding Table and is not found (see "Type Forwarding Tables" in "Forwarding Tables", p. 100). The Type Forwarding Table default action is set to pass so the packet is not discarded.
5. The packet is forwarded on WAN link 1 as the source MAC address is not allowed to forward packets on WAN link 2 by an entry in the Source Forwarding Table.

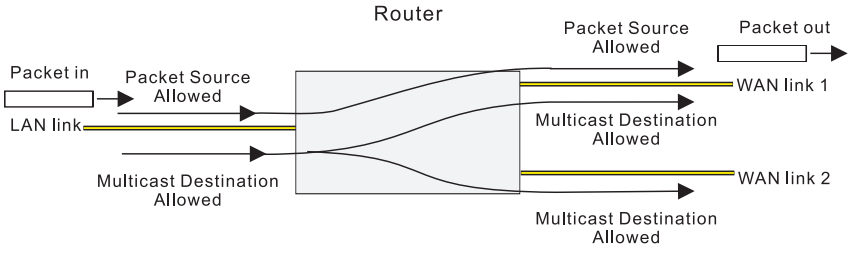

Controlled Bridging

Normal use of forwarding tables

Bridging services usually forward all or most packets received for bridging by default. Static forwarding tables can be set up to forward or discard specified packets (with a specified destination or source MAC address or packet type). That is, all packets are forwarded unless discarded by static forwarding table entries.

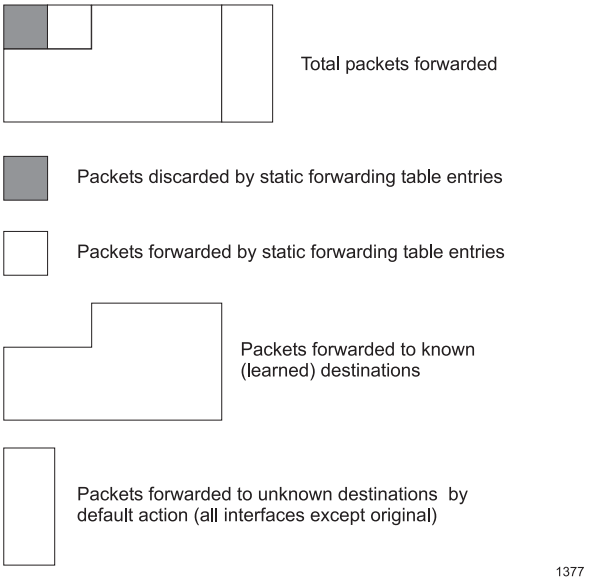
Total packets forwarded
Packets discarded by static forwarding table entries
Packets forwarded by static forwarding table entries
Packets forwarded to known (learned) destinations
Packets forwarded to unknown destinations by default action (all interfaces except original)
1377

The above diagram represents the data forwarded (unshaded area) of the total number of packets received for bridging.

Controlled bridging

The Intel Express 8100 Router allows packets which do not appear in the forwarding tables to be discarded. In addition, the learning function can be disabled so that the bridge forwarding tables only contains the defined static entries. When set up in this way, the bridge only forwards packets addressed to defined MAC addresses, from defined source addresses and with a specified Ethernet packet type. That is, all packets are discarded unless otherwise allowed to pass by static forwarding table entries.

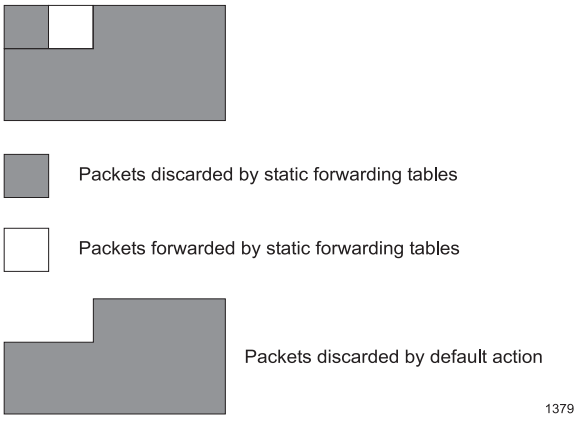
Packets discarded by static forwarding tables
Packets forwarded by static forwarding tables
Packets discarded by default action
1379

The above diagram represents the data forwarded (unshaded area) of the total number of packets received for bridging.

This way of controlling bridged data traffic can be used for security purposes and to restrict the bridged data traffic over expensive WAN links.

Controlled bridging with bridge learning

It is still possible to have controlled bridging with bridge learning enabled. In this case, the bridge still forwards packets to learned and defined (static) destinations if permitted by the Source and Type Forwarding Tables. That is, only packets with unknown destinations or discarded by static Destination Forwarding Table entries are discarded by the Destination Forwarding Tables.

Packets discarded by static forwarding tables
Packets forwarded by static forwarding tables
Packets forwarded to known (learned) destinations
Packets discarded by default action to unknown destinations
1378

The above diagram represents the data forwarded (unshaded area) of the total number of packets received for bridging.

Note One of the routers in a bridges over a WAN link must forward packets by default. Otherwise the other router cannot learn remote destinations.

Spanning Tree Algorithm

Bridging Loops

Problem with conventional bridges

The following figure shows a simple Local Area Network built up of 3 segments connected via 2 bridges.

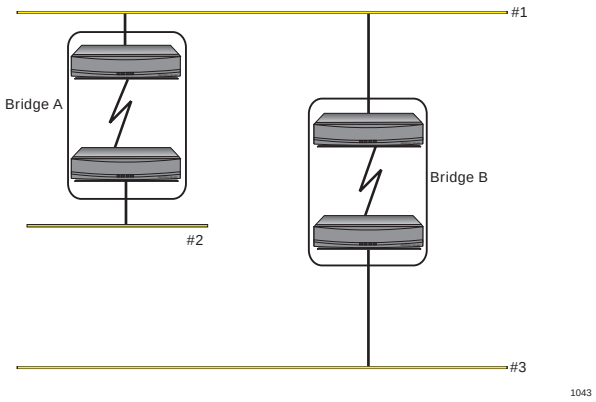

Note In the case of the Intel Express 8100 Router, a bridge between two LAN segments consists of two routers connected via a WAN link.

If either of the two bridges fails, one of the network segments is isolated. To reduce the possibility of one of the bridges becoming isolated, it is common practice to add a third bridge—if any one of the bridges fails, all three network segments are still accessible via the other two bridges.

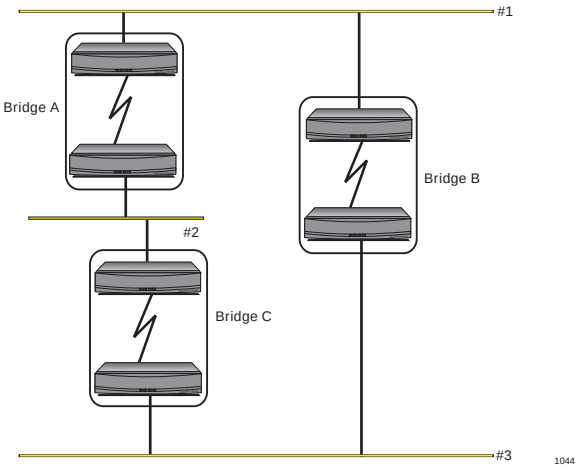

Using basic bridges which do not implement the Spanning Tree Algorithm, this configuration cannot be implemented—the bridges transfers received packets on one link to the network segments connected to the other links. This means that packets can continue to cycle within the network segments indefinitely—this is known as a "bridging loop". The receiver receives many copies of the same packet.

A solution to this problem is described in the following section.

Spanning Tree Protocol

Since alternative paths are desirable for backup and other purposes, IEEE and ISO have drawn up a proposal for a standard which offers a solution to "network loops" problem. The solution is called "The Spanning Tree Protocol", and is described in IEEE document 802.1D "Local MAC Bridges".

Spanning Tree Concept

Basic requirements

The following gives an overview of the basic requirements of the Spanning Tree Algorithm for bridges. For full details, see IEEE 802.1D.

Bridging loop detection and removal

Within the Spanning Tree Algorithm, bridges connected in a LAN must be capable of detecting potential “bridge loops”, and of removing these loops as they occur by blocking the appropriate links to other bridges. This is illustrated in the following diagram:

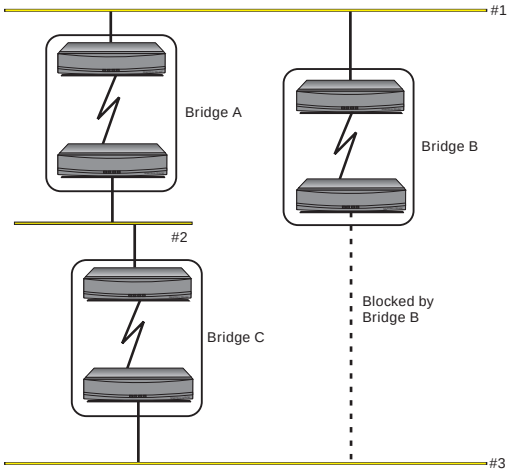

An alternate path has been established by connecting Bridge B in parallel with Bridges A and C, which also creates a potential bridging loop. However, by using the Spanning Tree Algorithm, Bridge B breaks the loop by blocking its path to segment 3.

Bridge failure

In the event that one of the bridges fails, the Spanning Tree Algorithm must be capable of activating alternative paths.

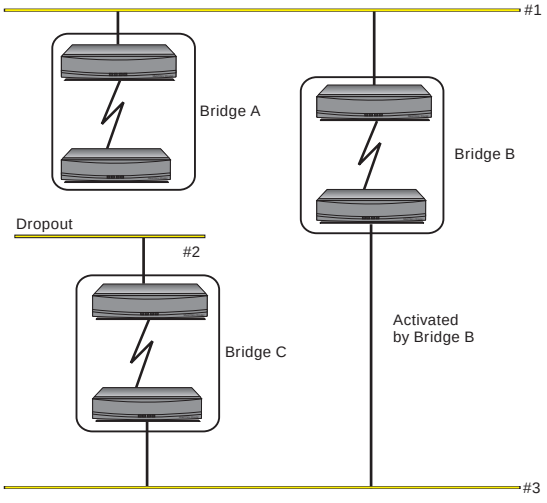

Network extension

If the network is extended, the Spanning Tree Algorithm must be capable of adapting automatically to the new topology.

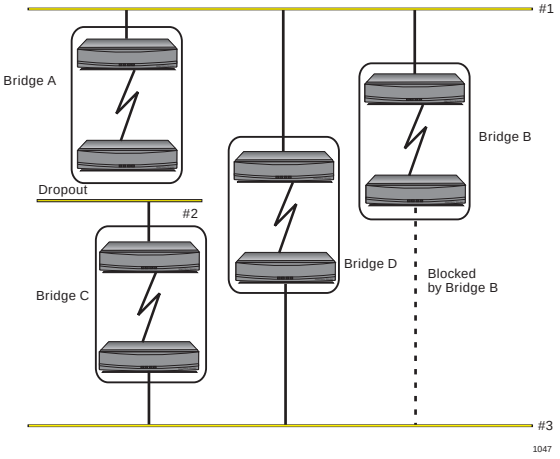

Spanning Tree in the Intel Express 8100 Router

Bridges and links

Each bridge is identified by a Bridge ID, and each link on a bridge is identified by a Port ID.
Bridge links can be either disabled or enabled. links which are enabled can be in one of the following states:

Listening

The initial state for a bridge, in which the bridges send messages to one another to establish the network topology and the optimal paths to the different segments of the network. Normal data is not transmitted.

Blocking

The bridge enters the Blocking State if a path with higher priority is found to exist during the Listening State. Normal data is not transmitted.

Learning

The bridge enters the Learning State if no path with a higher priority is found during the Listening State. Learned entries are entered in the Unicast Destination Forwarding Table. Normal data is not transmitted.

Forwarding

The bridge enters the Forwarding State after having been in the Learning State for a predefined time period. Normal data is transmitted.

Bridge ports and router links

In bridging, “port” is used to describe each connection to the bridge. In the router, “port” used to describe the physical LAN and WAN ports, and WAN and ports can support more than one WAN link. For the purposes of describing bridging correctly, the words “ports” and “links” are interchangeable when describing bridge ports within the router.

Root Bridge

During the Listening State, the bridge with the highest priority is determined. This bridge is then defined as the “Root Bridge”.

Priorities and cost factors

A priority is defined for each bridge and to each link of the bridge, the priority for the bridge is defined by the `Bridge Priority` parameter and the priority for a link, is defined by the `Port Priority` parameter. Each link also has a “cost factor” assigned to it which uses the Path Cost defined by the `Path Cost` parameter.

These parameters are used by the Spanning Tree Algorithm to determine the topology of the bridged network.

Spanning tree topology

The cost factor is used to calculate the distance from each link of a bridge to the Root Bridge. On the basis of this, each link on a bridge is assigned one of the following states:

Root Port

“Root Port” is assigned to links that is closest to the Root Bridge. Only one link on each bridge is assigned as the Root Port.

Designated Port

“Designated Port” is assigned to links that connect to bridges further away from the Root Bridge than the bridge. The Root Bridge can only have Designated and Blocking Ports.

Blocking

If any links are not assigned a Root Port or a Designated Port State, they are assigned a “Blocking” State. When a link is in the Blocking State, with the exception of Configuration BPDUs, packets are not accepted or transmitted by the link. The link can be said to be in standby.

Packet propagation

By enforcing this strict hierarchy, and by only forwarding packets between Root Ports and Designated Ports, the possibility of bridging loops is removed.

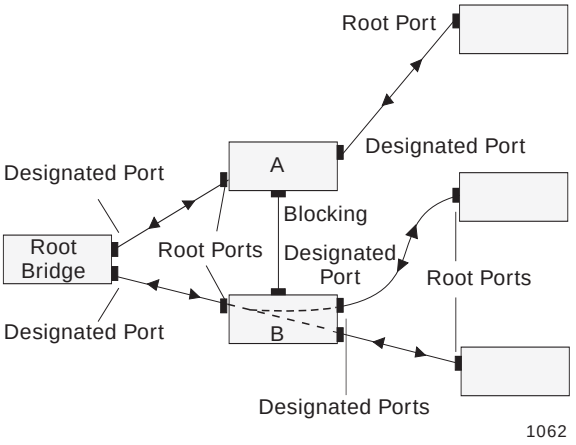

In the above diagram, packets sent from Bridge B to Bridge A are discarded by the blocking port on Bridge A. Packets are never forwarded from the blocking port on Bridge A. Packets can therefore not be sent directly between Bridges A and B except via the Root Bridge.

7 hop limit

In addition to the strict bridging hierarchy imposed by the Spanning Tree Algorithm, a 7 hop limit is introduced. This means that any packets cannot pass more than 7 bridges. This effectively limits the practical size of the bridged network when using the Spanning Tree Algorithm.

Configuration BPDU messages

To establish the optimal paths (during the Learning State), each bridge sends configuration BPDU packets to its neighboring bridges. These Configuration BPDU messages contain information about the spanning tree topology. The contents of these packets only changes when the bridged network topology changes or has not been established.

When the optimal paths within the network have been established, each bridge continues to send configuration BPDU packets to its neighboring bridges to tell them that it is still alive.

Configuration BPDU message propagation

When a bridged network is in a stable condition, bridges continue to send configuration BPDU packets to its neighboring bridges at regular intervals. Configuration BPDU packets are transmitted down the spanning tree from Designated Ports to Root Ports. If a Configuration BPDU is not received by the Root Port on a bridge in a predefined time interval (for example because a bridge along the path has dropped out), the link enters the Listening State to redetermine its optimal path.

Bridge Spoofing

BPDU Spoofing

Configuration BPDU packets over WAN links

Configuration BPDU packets used by the Spanning tree Algorithm take up valuable bandwidth on any type of WAN link. An example of this occurs when communicating over dial-up (switched) WAN links. Configuration BPDU packets are sent from Designated Ports to Root Ports at intervals defined by the Hello Time in the Root Bridge (2 seconds as the default). These configuration BPDU packets can prevent dial-up (switched) WAN links from becoming timecut (inactive)—the operating costs are then excessive.

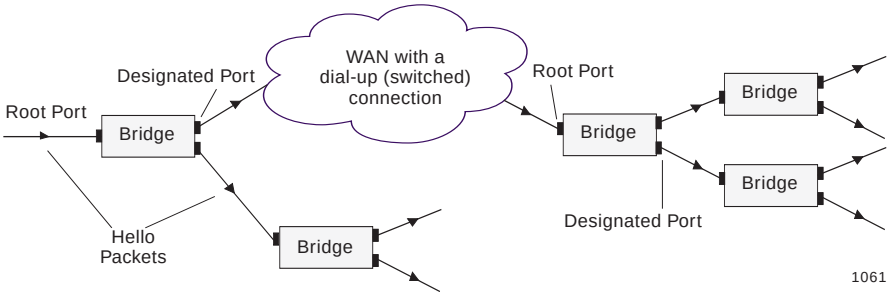

Bridge Spoofing

When the bridging topology is stable (when no bridges are being added or removed from the bridged network), the contents of Configuration BPDU packets sent between the bridges is the same. The contents of these packets only changes when the bridged network topology changes.

To counteract the problem of Configuration BPDU packets taking up excessive WAN link bandwidth the Intel Express 8100 Router implements Bridge Spoofing. When Bridge Spoofing is set up at the bridges at either side of a WAN link, the WAN link is not used to send Configuration BPDU packets from a Designated Port to a non-designated port unless the contents of the Configuration BPDU packet changes.

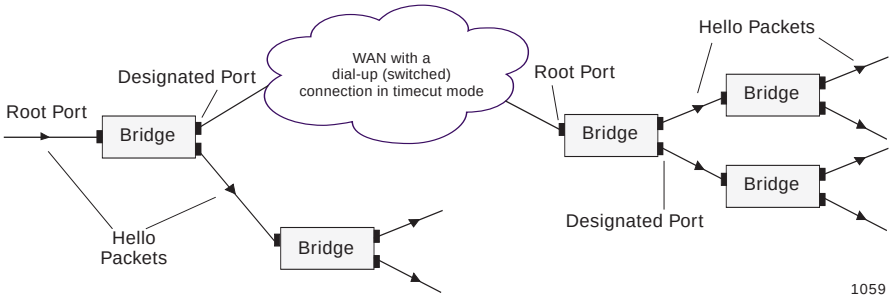

Bridge Spoofing on a Designated Port

When Bridge Spoofing is set up on a Designated Port, the link only forwards configuration BPDU packets when the contents of the configuration BPDU packets changes, that is, when the topology of the bridged network changes or is being established. When the contents of a configuration BPDU packet changes, configuration BPDU packets are sent over the WAN link until a valid acknowledge is received from the non-designated port on the bridge over the WAN link.

Bridge Spoofing on a non-designated port

When Bridge Spoofing is set up for a non-designated port, the bridge does not expect configuration BPDU packets from the non-designated port. The bridge continues to send the last configuration BPDU packet it received at the intervals defined by the Hello Time in the Root Bridge, to bridges connected on Designated Ports to maintain the integrity of the bridged network.

Non-BPDU Spoofing

Bridge Spoofing on a WAN link

Another feature of Bridge Spoofing is when two bridge ports are connected via a WAN link and one port is in the Blocking State and the other is in the Forwarding State. Normally, Non-BPDU packets are sent from the Forwarding Port and are Discarded by the Blocking Port—this is a waste of valuable WAN resources. When Bridge Spoofing is enabled on the bridges, no Non-BPDU packets (of any type) are forwarded to a Blocking Port.

Enabling Bridge Spoofing

When can Bridge Spoofing be used

Bridge Spoofing can only be used between two Intel Express 8100 Routers. Bridge Spoofing cannot be used on WAN links connected to routers/bridges from other manufacturers.

Requirements for use

Bridge Spoofing must be set up at both the routers on either side of a WAN link. for satisfactory operation.

Enabling Bridge Spoofing

Bridge Spoofing is set up when setting up bridging services for a WAN link, on the `Advanced` screen.

Note Bridge Spoofing is enabled by default for WAN links. When communicating with a bridge which does not support Bridge Spoofing, for example a bridge from another manufacturer, Bridge Spoofing must be disabled.

Appendices

Appendix A

Protocol Bandwidth Overheads and Requirements

Introduction

Dynamic network protocols (IP RIP and IPX RIP and SAP) and Spanning Tree bridging services use supervisory packets to communicate between network devices of, for example, the reachable networks and services and the network topology. These supervisory packets take link bandwidth which is particular problem over a WAN links which generally have a more restricted bandwidth. This appendix allows you to calculate the bandwidth requirements for your network.

Note Many facilities are built into the Intel Express 8100 Router to restrict these network overheads (e.g., triggered RIP updates, configurable IPX RIP and SAP updates over WAN links, IPX watchdog and SPX keep-alive spoofing and bridge BPDU spoofing). These are not accounted for in the calculation examples.

Bandwidths used by IP RIP

The following table lists the bandwidth used by IP RIP:

- IP RIP header included in every RIP packet takes 32 bytes
- IP network entries included in a RIP packet take 20 bytes/entry
- each IP RIP packet can only contain 25 entries, i.e., if more than 25 stations are transmitted, 2 or more packets must be used
- IP RIP updates are sent every 30 seconds

Example bandwidth usages for IP RIP

For 50 IP networks, the IP RIP protocol requires 284 bit/s of link bandwidth. For 100 IP networks, the IP RIP protocol requires 568 bit/s of link bandwidth, and so on.

Bandwidths used by IPX RIP and SAP

The following table lists the bandwidth used by IPX RIP and SAP:

- IPX RIP and SAP header included in every RIP and SAP packet takes 32 bytes
- IPX network entries included in a RIP packet take 8 bytes/entry
- IPX service entries included in a SAP packet take 64 bytes/entry

- each IPX RIP packet can only contain 50 entries, i.e., if more than 50 stations are transmitted, 2 or more packets must be used
- each IPX SAP packet can only contain 7 entries, i.e., if more than 7 stations are transmitted, 2 or more packets must be used
- IPX RIP and SAP updates are sent every 60 seconds

Example bandwidth usages for IPX RIP

For 50 IPX networks, the IPX RIP protocol requires 52 bit/s of link bandwidth. For 100 IPX networks, the IPX RIP protocol requires 115 bit/s of link bandwidth, and so on.

Example bandwidth usages for IPX SAP

For 50 IPX services, the IPX SAP protocol requires 460 bit/s of link bandwidth. For 100 IPX services, the IPX SAP protocol requires 917 bit/s of link bandwidth, and so on.

Bandwidths used by Spanning Tree bridging services

The following table lists the bandwidth used by Spanning Tree bridging:

- each bridging BPDU packet is 35 bytes/link
- bridging BPDU packets are sent every second

Example bandwidth usages for Spanning Tree bridging

Bridging updates (BPDU packets require 280 bit/s of link bandwidth.

Appendix B

TCP and UDP Port Numbers

Introduction

Ports are used by the Transport Control Protocol (TCP) and User Data Protocol (UDP) to name the ends of logical connections which carry long-term associations. For the purpose of providing services to unknown callers, a service contact port is defined. The following list specifies the ports used by the server process as its contact port. The contact port is sometimes called the 'well-known port'. The assigned ports use a small portion of the possible port numbers. The assigned ports have all bits set to zero except the low-order eight bits

Port assignments

The following table provides an extract of RFC 1700 of the most-used port numbers—low-order eight bits of the port are specified. Additional port numbers are listed in RFC 1700.

Decimal	Keyword	Protocol	Description
20	FTP-DATA	TCP	File Transfer [Default Data]
21	FTP	TCP	File Transfer [Control]
23	TELNET	TCP	Telnet
25	SMTP	TCP	Simple Mail Transfer Protocol
49	LOGIN	TCP	Login Host Protocol
53	DNS	TCP / UDP	Domain Name Service
63	VIA-FTP	TCP	VIA Systems—FTP
67	BOOTPS	UDP	Bootstrap Protocol Server
68	BOOTPC	UDP	Bootstrap Protocol Client
69	TFTP	UDP	Trivial File Transfer
70	Gopher	TCP	Gopher file service.
80	WWW	TCP	World Wide Web services.
161	SNMP	UDP	Simple Network Management Protocol.

Decimal	Keyword	Protocol	Description
162	SNMP Trap	UDP	Simple Network Management Protocol Traps.
520	RIP	UDP	Routing Information Protocol.
2049	NFS	UDP	Network File Servers.

Appendix C

Ethernet Type Codes

Ethernet Type Codes

The following table lists the known Ethernet type codes.

Type field	Description
0000-5DC	IEEE 802.3 length field
0101-01FF	Experimental—for development (conflicts with 802.3 length field)
0200	Xerox PUP conflicts with IEEE 802.3 length field
0201	Xerox PUP Address Translation (conflicts with IEEE 802.3 length field)
0600	XNS Internet
0800	DoD Internet
0801	X.75 Internet
0802	NBS Internet
0803	ECMA Internet
0804	CHAOSNet (proposed by Symbolic)
0805	X.25 Level 3
0806	Ethernet Address Resolution ARP (proposed by Symbolic)
0807	XNS Compatibility
081C	Symbolics Private
0888-088A	Xyplex
0900	Ungermann-Bass Network Debugger
0A00	Xerox IEEE 802.3 PUP
0A01	Xerox IEEE 802.3 PUP Address Translation
0BAD	Banyan VINES IP
0BAE	Banyan VINES Loopback

Type field	Description
0BAF	Banyan VINES Echo
1000	Berkeley trailer negotiation
1001-100F	Berkeley trailer encapsulation for IP
1600	VALID system protocol
4242	PCS Basic Block Protocol
5208	BBN Simnet Private
6000	DEC unassigned
6001	DEC Maintenance Operation Protocol (MOP) Dump/Load Assistance
6002	DEC MOP Remote Console
6003	DEC DECNet Phase IV
6004	DEC Local Area Transport (LAT)
6005	DEC DECNet Diagnostics
6006	DEC DECNet Customer Use
6007	DEC Local Area VAX Cluster (LAVC)
6008	DEC unassigned
6009	DEC unassigned
6010-6014	3Com Corporation
7000	Ungermann-Bass (UB) download
7001-7002	UB diagnostic/loopback
7020-7029	LRT (England)
7030	Proteon
8003	Cronus VIN
8004	Cronus Direct
8005	HP Probe protocol
8006	Nestar
8008	AT&T
8010	Excelan
8013	Silicon Graphics diagnostic (obsolete)
8014	Silicon Graphics network games (obsolete)

Type field	Description
8015	Silicon Graphics reserved type (obsolete)
8016	Silicon Graphics XNS NameServer, bounce server (obsolete)
8019	Apollo Domain
802E	Tymshare
802F	Tigan, Inc.
8035	Reverse Address Resolution Protocol RARP
8036	Aeonic Systems
8038	DEC LANBridge Management
8039	DEC unassigned
803A	DEC unassigned
803B	DEC unassigned
803C	DEC unassigned
803D	DEC Ethernet CSMA/CD Encryption Protocol
803E	DEC unassigned
803F	DEC LAN Traffic Monitor Protocol
8040	DEC unassigned
8041	DEC unassigned
8042	DEC unassigned
8044	Planning Research Corp.
8046-8047	AT&T
8049	ExperData (France)
805B	Versatile Message Translation Protocol RFC-1045 (Stanford)
805C	Stanford V Kernel, production
8060	Little Machines
8062	Counterpoint Computers
8065-8066	University of Massachusetts at Amherst
8067	Veeco Integrated Automation
8068	General Dynamics
8069	AT&T
806A	Autophon (Switzerland)

Type field	Description
806C	ComDesign
806D	Compugraphic Corporation
806E-8077	Landmark Graphics Corporation
807A	Matra (France)
807C	Merit Internodal
807D-8080	Vitalink Communications
8080	Vitalink TransLAN III Management
8081-8083	Counterpoint Computers
8088-808A	Xyplex
809B	Kinetics EtherTalk (AppleTalk over Ethernet)
809C-809E	Datability
809F	Spider Systems Ltd.
80A3	Nixdorf Computers (Germany)
80A4-80B3	Siemens Gammasonics Inc.
80C0-80C3	Digital Communications Assoc. Inc. DCA Data Exchange Cluster
80C4	Banyan VINES IP
80C5	Banyan VINES Echo
80C6	Pacer Software
80C7	Applitek Corporation
80C8-80CC	Intergraph Corporation
80CD-80CE	Harris Corporation
80CF-80D2	Taylor Instrument
80D3-80D4	Rosemount Corporation
80D5	IBM SNA Services over Ethernet
80DD	Varian Associated
80DE	Integrated Solutions Transparent Remote File System (TRFS)
80DF	Integrated Solutions
80E0-80E3	Allen-Bradley
80E4-80F0	Datability

Type field	Description
80F2	Retix
80F3	Kinetics AppleTalk Address Resolution Protocol (AARP)
80F4-80F5	Kinetics
80F7	Apollo Computer
80FF-8103	Wellfleet Communications
8069	AT&T
807B	Dansk Data Electronik A/S
8107	Symbolics Private
8108	Symbolics Private
8109	Symbolics Private
8130	Waterloo Microsystems Inc.
8131	VG Laboratory Systems
8137	Novell NetWare IPX (old)
8138	Novell, Inc.
8139-813D	KTI
9000	Loopback (Configuration Test Protocol)
9001	Bridge Communications XNS Systems Management
9002	Bridge Communications TCP/IP Systems Management
FF00	BBN VITAL LAN Bridge cache wakeups

Appendix D

IP Frame Formats

IP packets

The frame format of an IP packet is as follows:

IP header	
Version	4 bits
Internet Header Length (ILH)	4 bits
Type of service	1 byte
Total length	2 bytes
Identification	2 bytes
Flags	3 bits
Fragment offset	13 bits
Time to live	1 byte
Protocol	1 byte
Header checksum	2 bytes
Source address	4 bytes
Destination address	4 bytes
Options (if any)	-
Data	

IP Ethernet Type header

An IP Ethernet Type header looks like this:

Destination Address	6 bytes
Source Address	6 bytes
Frame Type	2 bytes = 0800

Appendix E

Novell IPX Frame Formats

Introduction

Novell IPX supports different frame types—the frame type supported depends on the Novell driver.

This appendix describes the different frame type supported by Novell IPX.

IPX packets

The frame format of an IPX packet is as follows:

IPX header (802.3, 802.2, SNAP, Ethernet Type)	
Checksum	2 bytes = FFFF
Packet Length	2 bytes
Transport Control	1 byte
Packet type	1 byte
Destination Network	4 bytes
Destination Node	6 bytes
Destination Socket	2 bytes
Source Network	4 bytes
Source Node	6 bytes
Source Socket	2 bytes
Data	
Pad	
Frame Check Sequence (FCS)	

IPX Ethernet 802.3 header

An IPX Ethernet 802.3 header looks like this:

Destination Address	6 bytes
Source Address	6 bytes
Length	2 bytes

IPX Ethernet 802.2 header

An IPX Ethernet 802.2 (also called LSAP) header looks like this:

Destination Address	6 bytes
Source Address	6 bytes
Length	2 bytes
DSAP	1 byte = E0
SSAP	1 byte = E0
Control field	1 byte

IPX SNAP header

An IPX SNAP header looks like this:

Destination Address	6 bytes
Source Address	6 bytes
Length	2 bytes
DSAP	1 byte = AA
SSAP	1 byte = AA
Control field	1 byte = 03
Protocol identification	5 bytes = 0000008137

IPX Ethernet Type header

An IPX Ethernet Type (also called Ethernet II) header looks like this:

Destination Address	6 bytes
Source Address	6 bytes
Frame type	2 bytes = 8137

Token Ring 802.2 header

The DSAP and SSAP contained in a Token Ring frame indicates that the frame contains an IPX packet. When Route fields are present, i.e. Route Control Broadcast etc., source route bridging is used. A Token Ring 802.2 header looks like this:

Access Control	1 byte
Frame Control	1 byte
Destination Address	6 bytes
Source Address	6 bytes
Route Control Broadcast / Length	1 byte

Route Control Direction	1 byte
Route Designator 1	2 bytes
....	2 bytes
Route Designator n	2 bytes
DSAP	1 byte = E0
SSAP	1 byte = E0
Control	1 byte = 03

Token Ring SNAP

The DSAP and SSAP contained in a Token Ring SNAP packet indicates SNAP encapsulation. Protocol identification indicates that the frame contains an IPX packet. When Route fields are present, i.e. Route Control Broadcast etc., source route bridging is used. A Token Ring 802.2 header looks like this:

Access Control	1 byte
Frame Control	1 byte
Destination Address	6 bytes
Source Address	6 bytes
Route Control Broadcast / Length	1 byte
Route Control Direction	1 byte
Route Designator 1	2 bytes
....	2 bytes
Route Designator n	2 bytes
DSAP	1 byte AA
SSAP	1 byte AA
Control	1 byte 03
Protocol Identification	5 bytes 0000008137

Appendix F

Novell IPX Service Types

Service and Object types

The following table indicates different kinds of Novell services and associated objects that are currently known:

Service	Object
User	1
User Group	2
Print Queue	3
File Server	4
Job Server	5
Gateway	6
Print Server	7
Archive Queue	8
Archive Server	9
Job Queue	A
Administration	B
NAS SNA Gateway	21
NACS	23
Remote Bridge Server 2	24
Bridge Server	26
TCP/IP Gateway	27
EICON x25 Gateway	29
Time Sync Server	2d
Archive Server SAP	2e
Advert Print Server	47

Service	Object
Btrieve VAP 5.0	4b
SQL VAP/NLM	4c
Xtree Nw Version	4d
BTRIEVE VAP 4.xx	50
Print Queue user	53
EICON Router	58
ARCserver 3.0	66
WANcopy utility	72
TES Netware VMS	7a
Emerald Backup	92
NetWare Access Server	98
Portable NetWare	9e
PowerChute	a1
Compaq IDA Status Monitor	ac
RSPX Server	107
CSA MUX	114
CSA LCA	115
CSA CM	116
CSA SMA	117
CSA DBA	118
CSA NMA	119
CSA SSA	11a
CSA STATUS	11b
CSA APPC	11e
SNA TEST	126
CSA TRACE	12a
Communication Exec.	130
NNS Domain	133
NNS profile	135
NSS Queue	137

Service	Object
Compaq SNMP Agent	174
HP Laserjet	30c
ARCserve 4.0	3c4
NW SQL	3de
Sitelock	4808
Intel Netport (Print Server)	8002
WordPerfect Network v.	8888
Wildcard	FFFF

Index

Numerics

7 hop limit (bridging) 113

A

Abbreviations iv
Acronyms iv
Address
  notation (Novell) 72
  resolution (IP) 51
  Resolution Protocol (ARP) 51

B

Backup
  WAN Links - considerations for timecut . . 5
  WAN Links - description 4
Bandwidth-on-Demand (PPP Multilink)
  operation 40
Basic Rate Access (BRA) 21
Basic Rate Interface, ISDN 20, 21
Blocking state in the Spanning Tree Algorithm 111
Bridge
  default action for unknown destinations . . 97
  default action for unknown Ethernet frame types 100
  default action from unknown source Multicasts 99
  learning function 97
  Spoofing - introduction to 94
  spoofing (bridging) 114
Bridging
  and routing comparison 94
  background knowledge required 93
  Configuration BDPU messages 113
  controlled 106
  Destination Forwarding Tables - description 96
  Destination Forwarding Tables - operation of 102
  disabling bridge learning 97
  forwarding examples 104
  forwarding process 101
  Forwarding Tables filtering function . . . 100
  implementation of the Spanning Tree Algorithm 111
  Loop-Back check 94, 101
  loops 108
  Novell IPX 94
  operation 101
  purpose of in the router 93
  Source Forwarding Table description . . . 99
  Spanning Tree Algorithm 108
  standards 93
  TCP/IP 94
  Type Forwarding Table - description . . . 100
  when are frames bridged 94

C

Call back (PPP) 10
Called Party
  Number, ISDN 25
  Subaddress, ISDN 25
Calling Party
  Number, ISDN 25
  Subaddress, ISDN 25
Challenge Handshake Authentication Protocol (CHAP - PPP) 12
Channel
  bundling, see Multilink PPP
  types, ISDN 21

Compression
Frame Relay 19
PPP 11
Concept of the Spanning Tree in Bridging 109
Configuration
BPDU Frame spoofing 114
BPDU messages (bridging) 113
Connection types, ISDN 21
Controlled bridging 106
with bridge learning 107
Cost factor (Novell IPX) 86

D

Daily Activity Limit, description 7
Data
Communication-terminating Equipment (DCE) - definition of 33
encryption, Frame Relay 19
encryption, PPP 11
Link Connection Identifier (DLCI) 18
Terminal Equipment (DTE) - definition of 33
transmission process (Novell) 75
Data compression
Frame Relay 19
general 4
PPP 11
Default
bridge action for Multicasts from unknown sources 99
bridge action for unknown destinations . . 97
bridge action for unknown Ethernet frame types 100
bridge action for unknown sources 99
Delay (ticks - Novell IPX)
description 73
manually calculating for static routes 74
Destination
address check (X.25) 38
address check, ISDN 25
Forwarding Tables (bridging) 96
Dial-up (switched) links 3
Daily Activity Limit 7
Timecut mode 4
Disabling
bridge learning 97
Documents
related v
Dotted decimal notation 48
Dynamic
mapping (NAT) 60

E

Encapsulation protocol used by Frame Relay 18
Encryption
Frame Relay 19
PPP 11
Ethernet
802.2 header format (Novell IPX) 128
802.3 header format (Novell IPX) 127
services 2
type codes 121
Type header format (IP) 126

F

Features of Frame Relay 17
Filters
Bridging filters 100
IP (Firewall protection) - description 52
RIP (Novell IPX) - description 86
SAP (Novell IPX) - description 87
Firewall Protection, see Filters IP
Forwarding
concept (bridging) 95
examples (bridging) 104
phase in the Spanning Tree Algorithm . . 111
Tables (bridging) - use of 95
Forwarding tables
discarding Multicast packets from unknown sources 99
Frame Relay
Addressing 18
data compression 19
data encryption 19
elements definition 18
encapsulation 18
introduction 17
PVC management 18
Frame types 48
IP packets 126
IPX packets 127
LSAP 128
Novell IPX SNAP 128
Token Ring SNAP 129

H

Hops (Novell IPX)
- description 73
- manually calculating for static routes 74

I

ICMP protocol 52
Implementation
- of the Spanning Tree Algorithm 111

Interface
- designations in the Spanning Tree Algorithm 112
- states in the Spanning Tree Algorithm . 111

Internet Control Message Protocol (ICMP) . . 52
- Ping function 52

Internet Tunnel 41
- and Internet access considerations 42
- configuration overview on X.25 links . . . 44
- encryption 42
- IP static host router to the remote router . 42
- numbered IP link 42
- security over X.25 links application 43

IP
- address assignment 47
- address classes 48
- address notation 48
- Address Resolution Protocol (ARP) 51
- address structure 48
- available addresses 49
- bridging 94
- dotted decimal address notation 48
- Ethernet Type header format 126
- Filter - example 55, 84
- filter criteria 54, 83
- filtering with NAT 59
- Filters - description 52
- Frame Formats 126
- Internet Control Message Protocol (ICMP) 52
- logging filtered packets 55
- Network Address Translation (NAT) 57
- network numbers 48
- packet formats 126
- packets filtered 53
- Ping - description 52
- port number assignments 119
- private networks 47
- RIP metric 66
- RIP-1 and RIP-2 64
- static routes 68
- subnets 49
- triggered RIP 67
- Tunnel, see Internet Tunnel

IP routing
- static host route for Internet Tunnels 42
- use of numbered links for Internet Tunnels 42

IPX
- see Novell IPX
- WAN 2 Protocol 76

ISDN
- Basic Rate Interface (BRI) 21
- Bearer services 22
- channel types 21
- Concept 21
- connection types 21
- Numbering and addressing 24
- permanent connections 28
- Primary Rate Access (PRA) 21
- protocol variant supported 29
- standards supported 28
- Supplementary services 23
- Tele services 22
- use of SPIDs 29
- User-network interface reference configurations 30
- what is ISDN? 21
- what ISDN offers 21

K

Keep-alive (Novell SPX)
- packet spoofing 81
- packets 80
- packets over WAN links 80

L

LAN services concept 2
Learning phase in the Spanning Tree Algorithm 111
Leased lines 3
Link
- Activity Alarm 7
- Control Protocol (LCP) - PPP 9

Links
- Frame Relay 18

Internet Tunnel 41
ISDN 21
maximum number of 3
Point-to-Point (PPP) 8
PPP Multilink 39
PVCs (X.25) 34
SVCs (X.25) 34
Listening phase in the Spanning Tree Algorithm 111
Logging
filtered packets 55
Logical channels, X.25 34
Loop-Back check (bridging) 94, 101
Looping packets (Novell IPX) 76
Loops in Bridging 108
LSAP frame 128

M

MAC addressing 93
MAC Bridging, see bridging
Master link (PPP Multilinks) 39
Metric in Novell IPX 73
Multilink PPP 11, 39
Slave links 39
Multiple backup links 6

N

NAT - see Network Address translation 57
Negotiation over Novell IPX WAN Links . . . 77
Network Address Translation (NAT)
and filtering 59
and RIP updates 59
and routing 59
examples of use 62
introduction 57
on the LAN interface 59
Static and Dynamic Mapping 58
static mapping to internal addresses 62
Network numbers (Novell IPX) 73
Non-BPDU spoofing (bridging) 115
Novell IPX 71
address notation 72
addressing 72
bridging 94
cost reductions using RIP filtering 92
cost reductions using SAP filtering 92
data transmission process 75
Ethernet 802.2 header format 128
Ethernet 802.3 header format 127
Ethernet type header format 128
IPX WAN 2 Protocol 76
looping packets 76
metrics 73
metrics - manually calculating 74
metrics example 75
network number assignments 73
packet formats 127
protocol overview 71
RIP (Routing Information Protocol) 85
SAP (Service Advertising Protocol) 87
SAP filtering 87
serialization packets 79
sockets reserved by Novell 73
SPX, see Novell SPX
static routes and services 88
watchdog spoofing 79
Novell IPX routing
example use of a static backup route 89
Novell SPX
description 79
keep-alive packets 80
Number of WAN links available in the router 3
Numbered
and unnumbered IPX WAN links 76

O

Overview of PPP Multilink configuration . . . 39

P

Packet Switched Data Networks (PSDNs) . . 32
Password
Authentication Protocol (PAP - PPP) 15
Peer Authentication
CHAP, PPP 12
PAP (PPP) 15
Permanent ISDN 28
Permanent ISDN connections 28
Permanent Virtual Circuits (PVCs - X.25) . . 34
Ping function 52
Port
designations in the Spanning Tree Algorithm 112

number assignments (TCP/IP) 119
states in the Spanning Tree Algorithm . 111
PPP
call back 10
Challenge Handshake Authentication Protocol (CHAP) 12
data compression 11
data encryption 11
introduction to 8
Link Control Protocol 9
Multilink, configuration overview 39
Multilinks 11
Multilinks, see Multilink PPP
operation of 9
over HDLC 9
over ISDN 9
Password Authentication Protocol (PAP) 15
protocols 9
Tunnel, see Internet Tunnel
Primary Rate Access, ISDN 21
Protocols which cannot be routed 94
Proxy Address Resolution (Proxy ARP) 52
Purpose of Bridging in the router 93
PVC management, Frame Relay 18

R

Referenced documents
data encryption 12
for Firewall Protection 54
Related documents v
RFCs
1356 encapsulation (X.25) 37
877 encapsulation (X.25) 37
list of supported v
RIP (IP) 64
and NAT 59
differences between RIP-1 and RIP-2 ... 64
metric 66
RIP-1 and RIP-2 implementation 65
RIP-1 and RIP-2 in a network environment 65
RIP-1 and RIP-2 standards 64
triggered 67
RIP (Novell IPX) 85
filtering 86
Rx filters 86
Tx filters 86
Routing
and Bridging comparison 94
table (Novell IPX) 85

S

SAP (Novell IPX) 72
agent 87
broadcasts 87
filtering 87
Rx filters 88
Server Information Table 87
Services & object types 130
Tx filters 88
Semi-Permanent ISDN 28
Serialization packets (Novell IPX) 79
Server Information Table (Novell IPX) 87
Service Advertising Protocol (SAP - Novell IPX) 87
Slave links (PPP Multilink) 39
SNAP header format 128
Source
address check (X.25) 38
address check, ISDN 26
Forwarding Table (bridging) - use of ... 103
Forwarding Table (bridging), description 99
Forwarding Table (bridging), discarding Multicasts 99
Spanning Tree Algorithm (bridging) 108
7 hop limit 113
port (interface) designations 112
port (interface) states 111
SPIDs - use of by Nation ISDN-1 and AT&T ISDN 29
Spoofing
Novell IPX watchdog packets 79
Novell SPX keep-alive packets 81
SPX (Novell IPX/SPX) see Novell SPX
Static
and learned entries in forwarding tables (bridging) 97
mapping (NAT) 60
NAT for individual IP addresses 60
operating mode (PPP Multilinks) 40
routes (IP) 68
routes (IP) - remote boot provider 69
routes and services (Novell IPX) 88

Routes and Static Backup Routes (Novell IPX) 89
Subnet Masks 50
Subnets (IP) 49
Switched Virtual Circuits (SVCs - X.25) 34

T

TCP/IP, see IP
Timecut
backup WAN link considerations 5
mode 4
Timer Profiles
description 6
example 7
important considerations 6
Time-ticks (Novell IPX) 73
Token Ring
802.2 header format 128
SNAP frames 129
Total
Daily Activity Limit, see Daily Activity Limit
number of WAN links available 3
Triggered RIP (IP) 67
Triggered updates (Novell IPX) 86
Tunnel, Internet 41
Type
codes for Ethernet frames 121
Forwarding Table (bridging) - use of . . 103
Forwarding Table (bridging) description 100

U

Unnumbered and numbered IPX WAN links 76
Updating periods for RIP (Novell IPX) 86

V

Virtual Circuit (X.25)
Permanent (PVC) 34
Switched (SVC) 34

W

WAN Bridging
definition of 93
see Bridging
WAN links
Backup links 4
Dial-up (switched) links 3
general facilities available 3
Leased lines 3
maximum number of 3
types supported 3
Watchdog (Novell IPX)
packets 77
packets over WAN links 78
spoofing 79
When are frames routed and when are they bridged 94

X

X.121 numbering plan (X.25) 37
X.25
adding encryption and authentication 43
addressing 37
data security over using Internet Tunnels 43
device characterisation 33
encapsulation 37
general specifications 32
in the MatchBox Router 32
introduction to 32
introduction to Switched Data Networks (PSDNs) 32
line speed matching 34
logical channels 34
packet and window sizes 36
Permanent Virtual Circuits (PVCs) - description of 34
service types available 34
supervisory packets 35
Switched Virtual Circuits (SVCs) - description of 34
Virtual Call process 35
window sizes 36